

Phase diagram of matrix compressed sensing

Christophe Schülke*

*Laboratoire de Physique Statistique, CNRS, PSL Universités et Ecole Normale Supérieure, 75005, Paris, France and
Institut de Physique Théorique, CNRS, CEA, Université Paris-Saclay, 91191, Gif-sur-Yvette, France*

Philip Schniter†

*Department of Electrical and Computer Engineering,
The Ohio State University, Columbus, OH 43210, USA*

Lenka Zdeborová‡

Institut de Physique Théorique, CNRS, CEA, Université Paris-Saclay, 91191, Gif-sur-Yvette, France

In the problem of matrix compressed sensing we aim to recover a low-rank matrix from few of its element-wise linear projections. In this contribution we analyze the asymptotic performance of a Bayes-optimal inference procedure for a model where the matrix to be recovered is a product of random matrices. The results that we obtain using the replica method describe the state evolution of the recently introduced P-BiG-AMP [19] algorithm. We show the existence of different types of phase transitions, their implications for the solvability of the problem, and we compare the results of the theoretical analysis to the performance reached by P-BiG-AMP. Remarkably the asymptotic replica equations for matrix compressed sensing are the same as those for a related but formally different problem of matrix factorization [10].

* christophe.schulke@espci.fr

† schniter@ece.osu.edu

‡ lenka.zdeborova@cea.fr

I. INTRODUCTION

Recovering a sparse or a low-rank signal from as few observations as possible is a class of problems that attracted considerable attention in statistics and signal processing. Very popular examples of problems belonging to this class are compressed sensing [4], or matrix completion [3]. Another interesting member of this class is the problem matrix compressed sensing, in which one aims to recover a low-rank matrix from a few of its random component-wise linear projections. We give a formal definition of the problem in Sec. I A. This problem has a range of interesting applications, see [19] and references therein.

The main line of theoretical work related to matrix compressed sensing minimizes the nuclear norm of the matrix (i.e. the sum of its singular values) subject to the constraint that its linear projections agree with the measured values [5, 23]. Nuclear norm minimization is algorithmically tractable and when analyzed it provably recovers the unknown matrix for an interesting range of parameters. The nuclear norm is a common type of regularization that enhances low-rank solutions. A rank R matrix \mathbf{X} of dimension $M \times P$ can be written as a product of two matrices $\mathbf{X} = \mathbf{U}\mathbf{V}^\top$ of sizes $M \times R$ and $R \times P$. However, the nuclear norm minimization approach does not handle straightforwardly cases when there are some requirements (such as sparsity) on the factors \mathbf{U} and \mathbf{V} .

In the present paper we study the generalized matrix compressed sensing problem, where arbitrary component-wise constraints are put on the factors \mathbf{U} and \mathbf{V} and each of the linear projections is observed through some non-linear output channel. We are able to do that if we restrict to a probabilistic model where the components of the ground-truth factors \mathbf{U} and \mathbf{V} are i.i.d. random variables of known probability distribution, and where the probabilistic nature of the output channel is known. Under such assumptions the model is amenable to exact analysis via the replica method developed in statistical physics [17, 18]. The results stemming from the replica method are in general in one-to-one correspondence with the analysis of message passing algorithms designed to solve the problem in an optimal way, as illustrated for the compressed sensing problem in [11] or for matrix factorization in [10]. For the matrix compressed sensing problem this algorithm, called P-BiG-AMP, was derived and tested recently in [19]. Our contribution can hence also be viewed as an asymptotic analysis of the performance of this algorithm for the assumed model. We compare the analysis to the performance of P-BiG-AMP and indeed observe excellent agreement.

Our analysis reveals a striking connection between the matrix compressed sensing problem and the problem of matrix factorization as studied in [10, 12, 25]. These are two different inference problems. In matrix compressed sensing we observe a set of element-wise linear projections of the matrix, whereas in matrix factorization we observe the elements of the matrix directly. Yet the replica analysis of the two problems yields equivalent equations and hence the asymptotic behavior of the two problems, including the phase transition, is closely linked. This analogy was already remarked for the nuclear norm minimization for matrix compressed sensing and matrix denoising in [5], or for matrix compressed sensing and matrix completion in [24].

A. Definition of the problem

Let $\mathbf{X} \in \mathbb{R}^{M \times P}$ be a matrix of low rank $R < \min(M, P)$. It can thus be written as a product of two smaller matrices: $\mathbf{U} \in \mathbb{R}^{M \times R}$ and $\mathbf{V} \in \mathbb{R}^{P \times R}$,

$$\mathbf{X} = \mathbf{U}\mathbf{V}^\top. \quad (1)$$

The low-rank matrix compressed sensing problem consists in recovering \mathbf{X} from a set of linear combinations of its entries. We call $\mathcal{A} : \mathbb{R}^{M \times P} \rightarrow \mathbb{R}^L$ the linear operator associated to the matrix \mathbf{A} , we note

$$\mathbf{Z} = \mathcal{A}(\mathbf{X}) \in \mathbb{R}^L \quad (2)$$

and \mathbf{Y} the measured version of \mathbf{Z} after passing through an element-wise measurement channel:

$$\mathbf{Y} \sim p_Y^0(\mathbf{Y}|\mathbf{Z}). \quad (3)$$

This setting is shown in Fig. 1, and the goal is to reconstruct \mathbf{U} and \mathbf{V} (but sometimes only \mathbf{X}) from the knowledge of \mathbf{Y} .

We can rewrite (2) in the component-wise manner

$$\forall l \in [1, L], \quad z_l = \sum_{\mu=1}^M \sum_{p=1}^P A_l^{\mu p} x_{\mu p}. \quad (4)$$

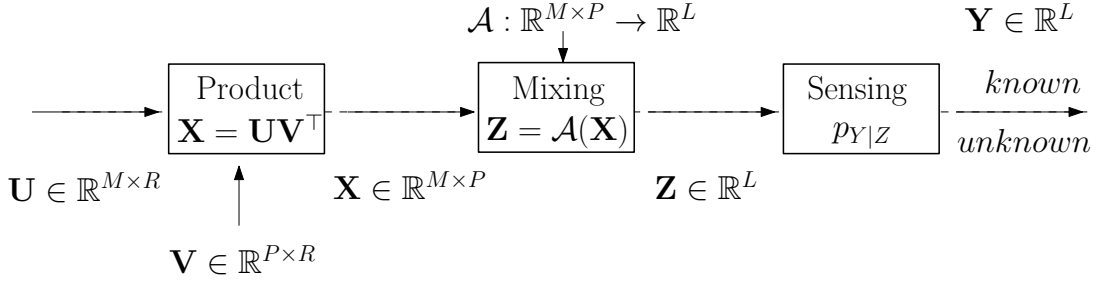


FIG. 1. The setting of generalized matrix compressed sensing. A low-rank matrix \mathbf{X} can be decomposed into a product of two smaller matrices \mathbf{U} and \mathbf{V} . A linear operator \mathcal{A} is applied to \mathbf{X} , producing an intermediary variable \mathbf{Z} . A measurement \mathbf{Y} of \mathbf{Z} is obtained through a noisy channel. The problem is closely linked to other inference problems: dropping the “mixing” block, one recovers a generalized matrix factorization problem. Dropping the “product” block, one recovers a generalized linear model.

1. The probabilistic model and assumptions of our analysis.

In order to enable the asymptotic analysis (i.e. when $M, P, L \rightarrow \infty$) via the replica method we introduce the following probabilistic model for matrix compressed sensing.

- We assume that elements of \mathbf{U} and \mathbf{V} are sampled independently at random such that

$$\mathbf{U} \sim \prod_{\mu s} p_U^0(u_{\mu s}), \quad \mathbf{V} \sim \prod_{ps} p_V^0(v_{ps}). \quad (5)$$

We assume the distributions p_U^0 and p_V^0 to have zero mean and respective variances Q_u^0 and Q_v^0 of order one. These distributions might not be known exactly: instead, we use zero-mean priors p_U and p_V believed to be close to p_U^0 and p_V^0 (in terms of Kullback-Leibler divergence).

- We assume the output distribution $p_{Y|Z}^0$ to be separable

$$p_{Y|Z}^0 = \prod_l p_{Y|Z}^0(y_l | z_l). \quad (6)$$

In the inference we use a distribution $p_{Y|Z}$ we believe to be close to $p_{Y|Z}^0$ (in terms of Kullback-Leibler divergence).

- We assume the matrix \mathbf{A} of the linear operator \mathcal{A} to have normally distributed i.i.d. elements with zero mean and variance $1/(RMP)$, such that the elements of \mathbf{Z} have zero mean and variance $Q_u^0 Q_v^0$. This is the same assumption as is often made in compressed sensing, and differentiates the problem from matrix factorization, in which \mathcal{A} is the identity.
- We assume the dimensions M , P and L to be large, but their following ratios to be of order one

$$\alpha_U = \frac{L}{RM}, \quad \alpha_V = \frac{L}{RP}. \quad (7)$$

On the other hand, R can be small.

2. Measures of recovery

Given the estimates $(\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{X}})$ that an algorithm returns for $(\mathbf{U}, \mathbf{V}, \mathbf{X})$, the following mean squared errors quantify how close the estimates are from the real values:

$$\text{MSE}_u = \frac{\|\mathbf{U} - \hat{\mathbf{U}}\|_F^2}{MR}, \quad \text{MSE}_v = \frac{\|\mathbf{V} - \hat{\mathbf{V}}\|_F^2}{PR}, \quad \text{MSE}_x = \frac{\|\mathbf{X} - \hat{\mathbf{X}}\|_F^2}{LR}, \quad (8)$$

where $\|\cdot\|_F$ is the Frobenius norm of a matrix. Note that as in matrix factorization, there is an inherent ill-posedness when it comes to recovering the couple (\mathbf{U}, \mathbf{V}) . As a matter of fact, for any $R \times R$ invertible matrix \mathbf{C} , the couple

$(\mathbf{U}\mathbf{C}, \mathbf{V}(\mathbf{C}^{-1})^\top)$ generates the same \mathbf{X} as (\mathbf{U}, \mathbf{V}) . In some case, this ill-posedness can be lifted thanks to the distributions p_U^0 and p_V^0 , but this is not always the case and might nevertheless be cause of trouble. In that case, it is possible to have a very low MSE_x but high MSE_u and MSE_v .

In the setting where $R = 1$, \mathbf{U} and $\hat{\mathbf{U}}$ are vectors and we can consider the following definitions of normalized mean squared errors

$$\text{nMSE}_u = 1 - \frac{|\mathbf{U}^\top \hat{\mathbf{U}}|}{\|\mathbf{U}\|_2 \|\hat{\mathbf{U}}\|_2}, \quad \text{nMSE}_v = 1 - \frac{|\mathbf{V}^\top \hat{\mathbf{V}}|}{\|\mathbf{V}\|_2 \|\hat{\mathbf{V}}\|_2} \quad (9)$$

that take values between 0 and 1 and take into account all invariances of the problem: an nMSE of 0 indicates perfect reconstruction up to the scaling invariance.

B. Notations

We use bold letters for vectors and matrices and non-bold letters for scalars. The elements of a vector \mathbf{x} are noted $[\mathbf{x}]_i$ or x_i . The operator \odot is used for element-wise multiplication of vectors or matrices. \mathbf{x}^{-1} , \mathbf{x}^2 and \mathbf{x}^\top refer respectively to the component-wise inverse, the component-wise square and the transpose of the vector (or matrix) \mathbf{x} . If \mathcal{A} is a linear operator and \mathbf{A} its matrix, we write \mathcal{A}^2 for the linear operator associated to \mathbf{A}^2 . We use the notation $\iota \equiv \sqrt{-1}$. Estimators \hat{X} and \hat{x} of a variable X are the minimal mean squared error (MMSE) estimators of estimated probability distribution functions $\hat{P}(x)$ and $\hat{p}(x)$. We note \bar{X} and \bar{x} the variances of these distributions and refer to them as *uncertainties*, as they are a measure of the uncertainty of the estimators \hat{X} and \hat{x} .

Using the matrix \mathbf{A} , we can define two auxiliary linear operators $\mathcal{A}_U : \mathbb{R}^P \rightarrow \mathbb{R}^{L \times M}$ and $\mathcal{A}_V : \mathbb{R}^M \rightarrow \mathbb{R}^{L \times P}$ such that

$$[\mathcal{A}_U(\mathbf{v})]_{l\mu} \equiv \sum_p A_l^{\mu p} v_p, \quad (10)$$

$$[\mathcal{A}_V(\mathbf{u})]_{lp} \equiv \sum_\mu A_l^{\mu p} u_\mu. \quad (11)$$

We note $x \sim p_X(x)$ a random variable x following the probability distribution p_X . This holds also for vectors and matrices: $\mathbf{x} \sim p_X(\mathbf{x})$. In that case, we say that $p_X(\mathbf{x})$ is separable if each component x_i of \mathbf{x} is sampled independently from the others: $\forall i, x_i \sim p_{X_i}(x_i)$, which we will note p_X as well if the components are identically distributed.

We write $f(x) \propto g(x)$ when the functions f and g are equal up to a multiplying constant that does not depend on x . We write $K = O(1)$ (respectively $K = O(M)$) in order to signify that K is of order 1 (respectively M).

Let us introduce some useful functions that will be used throughout the paper. We note $\mathcal{N}(x; \hat{x}, \bar{x})$ the normalized Gaussian with mean \hat{x} and variance \bar{x} :

$$\mathcal{N}(x; \hat{x}, \bar{x}) = \frac{1}{\sqrt{2\pi\bar{x}}} e^{-\frac{(x-\hat{x})^2}{2\bar{x}}}. \quad (12)$$

In integrals, we note Dt the integration over a variable t with a standard normal distribution:

$$Dt = dt \mathcal{N}(t; 0, 1). \quad (13)$$

For any function h and integer i , we define the i -th moment of the product of h multiplied by a Gaussian:

$$f_i^h(\hat{x}, \bar{x}) = \int dx x^i h(x) \mathcal{N}(x; \hat{x}, \bar{x}). \quad (14)$$

With (14), we define the mean and the variance of the distribution $\frac{h(x)\mathcal{N}(x; \hat{x}, \bar{x})}{f_0^h(\hat{x}, \bar{x})}$:

$$\hat{f}^h(\hat{x}, \bar{x}) = \frac{f_1^h(\hat{x}, \bar{x})}{f_0^h(\hat{x}, \bar{x})}, \quad (15)$$

$$\bar{f}^h(\hat{x}, \bar{x}) = \frac{f_2^h(\hat{x}, \bar{x})}{f_0^h(\hat{x}, \bar{x})} - \hat{f}^h(\hat{x}, \bar{x})^2, \quad (16)$$

It can be verified that following relations hold:

$$\frac{\partial}{\partial \bar{x}} f_i^h(\hat{x}, \bar{x}) = \frac{1}{\bar{x}} (f_{i+1}^h(\hat{x}, \bar{x}) - \hat{x} f_i^h(\hat{x}, \bar{x})), \quad (17)$$

$$\frac{\partial}{\partial \bar{x}} f_i^h(\hat{x}, \bar{x}) = \frac{1}{2\bar{x}^2} (f_{i+2}^h(\hat{x}, \bar{x}) - 2\hat{x} f_{i+1}^h(\hat{x}, \bar{x}) - (\bar{x} - \hat{x}^2) f_i^h(\hat{x}, \bar{x})), \quad (18)$$

$$\frac{\partial}{\partial s} f_i^h(\sqrt{st}, \rho - s) = -\frac{e^{\frac{t^2}{2}}}{2s} \frac{\partial}{\partial t} \left(e^{-\frac{t^2}{2}} \frac{\partial}{\partial t} f_i^h(\sqrt{st}, \rho - s) \right). \quad (19)$$

Finally, we introduce two further useful auxiliary functions:

$$\hat{g}^h(\hat{x}, \bar{x}) = \frac{\hat{f}^h(\hat{x}, \bar{x}) - \hat{x}}{\bar{x}}, \quad \bar{g}^h(\hat{x}, \bar{x}) = \frac{\bar{f}^h(\hat{x}, \bar{x}) - \bar{x}}{\bar{x}^2}. \quad (20)$$

II. ALGORITHMS

A. Message-passing algorithm

In this paper, we will focus on an approximate message passing (AMP) algorithm. AMP algorithms originated in studies of problems related to linear estimation [6, 9, 22]. For the above probabilistic model of matrix compressed sensing, AMP was derived and called P-BiG-AMP in [19]. In the following, we explain its principle and expose the main steps of its derivation.

In Bayesian inference, one seeks to produce estimators $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ of \mathbf{U} and \mathbf{V} using the following posterior probability:

$$p(\mathbf{U}, \mathbf{V} | \mathbf{Y}, \mathcal{A}) \propto p_U(\mathbf{U}) p_V(\mathbf{V}) p_{Y|Z}(\mathbf{Y} | \mathcal{A}(\mathbf{U}\mathbf{V}^\top)). \quad (21)$$

As explained above, the probability distributions used in (21) ideally match the distributions (3, 5) used for the generation of the problem, in which case the inference is said to be Bayes-optimal. However, it is often the case that these distributions are not known exactly: in this case, the distributions used in (21) are assumptions that we make on the signals' distributions and on the measurement channel. Inference is in that case suboptimal. However, in similar problems it has turned out that the results can still be satisfying despite the mismatch between the priors and the actual probability distributions. Furthermore, it is possible to parametrize the priors and learn the parameters during inference, for example using an expectation maximization procedure [1], which has proven to give satisfying results [11].

Starting from the posterior probability distribution (21), the two interesting questions are how to evaluate this quantity and how to obtain estimators $(\hat{\mathbf{U}}, \hat{\mathbf{V}})$ from it. For the second point, we will use the minimal mean squared error (MMSE) estimator, as our goal is to obtain low MSEs for (8). Concerning the first point, the problem in estimating (21) is that it is a distribution in a high-dimensional space. Though it is possible to sample from such a distribution using a Monte Carlo Markov chain, the procedure is very time consuming. Therefore we resort to loopy belief propagation (BP) to estimate the marginals of (21). Though not guaranteed to converge on this type of problems, BP has proven to be very successful in a variety of similar inference problems [11, 16].

In order to derive the BP algorithm, we first rewrite (21) to make all variables appear individually:

$$p(\mathbf{U}, \mathbf{V} | \mathbf{Y}, \mathcal{A}) \propto \prod_{\mu s} p_U(u_{\mu s}) \prod_{ps} p_V(v_{ps}) \int \prod_l dz_l p_{Y|Z}(y_l | z_l) \delta \left(z_l - \sum_{p=1}^P \sum_{\mu=1}^M A_l^{p\mu} \sum_{s=1}^R u_{ps} v_{\mu s} \right). \quad (22)$$

This probability distribution can be represented by the factor graph in figure 2. On it, two types of message pairs (m, \tilde{m}) and (n, \tilde{n}) are sent to and from the u and v variables respectively. As the roles of u and v are completely symmetric, we will only treat explicitly the pair (m, \tilde{m}) : the result can be generalized straightforwardly to (n, \tilde{n}) . The message-passing update equations read:

$$m_{\mu s \rightarrow l}^{t+1}(u_{\mu s}) \propto p_U(u_{\mu s}) \prod_{l' \neq l} \tilde{m}_{l' \rightarrow \mu s}^t(u_{\mu s}), \quad (23)$$

$$\tilde{m}_{l \rightarrow \mu s}^{t+1}(u_{\mu s}) \propto \int \left(\prod_{ps'} dv_{ps'} n_{ps' \rightarrow l}^{t+1}(v_{ps'}) \prod_{(s', \mu') \neq (s, \mu)} du_{\mu s'} m_{\mu s' \rightarrow l}^{t+1}(u_{\mu s'}) \right) dz p_{Y|Z}(y_l | z) \delta(z - \mathcal{A}(\mathbf{U}\mathbf{V}^\top)), \quad (24)$$

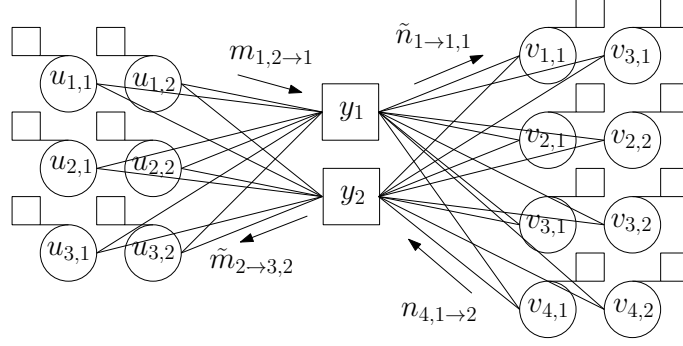


FIG. 2. Factor graph associated to the probability distribution (22). Here, we used $R = 2$, $M = 3$, $P = 4$, $L = 2$. Circle represent variables, squares represent constraints. The small squares represent the priors on the variables u and v . Messages $(m, \tilde{m}, n, \tilde{n})$ are sent along each edge of the factor graph.

where the \propto sign stands because (m, \tilde{m}) are probability distributions and must therefore be normalized. These equations can be seen as fixed point equations or as iterative equations that constitute an algorithm. For notational lightness, we will do the following calculations without time indices. However, the correct time indices are crucial for the final algorithm to converge.

A first simplification can be made by replacing the $R(M + P)$ integrals in (24) by a single one over the variable z , which is the sum of $R(M + P) - 1$ random variables. In BP, we assume these random variables to be independent, which allows us to use the central limit theorem. Calling $\hat{u}_{\mu s \rightarrow l}$ and $\bar{u}_{\mu s \rightarrow l}$ respectively the means and variances of the variable $u_{\mu s}$ distributed according to the distribution $m_{\mu s \rightarrow l}$ (and similarly for the variables v_{ps}), the variable $z_l = \sum_{\mu p} A_l^{\mu p} \sum_s u_{\mu s} v_{ps}$ is a Gaussian variable with mean and variance:

$$\hat{Z}_l = \sum_{\mu p s} A_l^{\mu p} \hat{u}_{\mu s \rightarrow l} \hat{v}_{ps \rightarrow l}, \quad (25)$$

$$\begin{aligned} \bar{Z}_l &= \sum_{\mu p s} (A_l^{\mu p})^2 [\bar{u}_{\mu s \rightarrow l} \bar{v}_{ps \rightarrow l} + (\hat{u}_{\mu s \rightarrow l})^2 \bar{v}_{ps \rightarrow l} + \bar{u}_{\mu s \rightarrow l} (\hat{v}_{ps \rightarrow l})^2] \\ &+ \sum_{ps} \sum_{\mu \neq \mu'} A_l^{\mu p} A_l^{\mu' p} \bar{v}_{ps \rightarrow l} \hat{u}_{\mu s \rightarrow l} \hat{u}_{\mu' s \rightarrow l} \\ &+ \sum_{\mu s} \sum_{p \neq p'} A_l^{\mu p} A_l^{\mu p'} \bar{u}_{\mu s \rightarrow l} \hat{v}_{ps \rightarrow l} \hat{v}_{p' s \rightarrow l}. \end{aligned} \quad (26)$$

However, in eq. (24), $u_{\mu s}$ is fixed and thus $(\hat{u}_{\mu s \rightarrow l}, \bar{u}_{\mu s \rightarrow l})$ has to be replaced by $(u_{\mu s}, 0)$ in (25,26). Defining $(\hat{Z}_{l \rightarrow \mu s}, \bar{Z}_{l \rightarrow \mu s})$ to be (\hat{Z}_l, \bar{Z}_l) with $(\hat{u}_{\mu s \rightarrow l}, \bar{u}_{\mu s \rightarrow l}) = (0, 0)$ and

$$F_{l\mu s} = \sum_p A_l^{\mu p} \hat{v}_{ps \rightarrow l}, \quad (27)$$

$$H_{l\mu s} = 2 \sum_p \sum_{\mu' \neq \mu} A_l^{\mu p} A_l^{\mu' p} \hat{u}_{\mu' s \rightarrow l} \bar{v}_{ps \rightarrow l}, \quad (28)$$

$$G_{l\mu s} = \sum_p (A_l^{\mu p})^2 \bar{v}_{ps \rightarrow l}, \quad (29)$$

one can rewrite (24) with a single integral over a variable z following a Gaussian distribution. Using the definition (14), the message (24) can be expressed as a simple function of the mean and variance of this Gaussian:

$$\tilde{m}_{l \rightarrow \mu s}(u_{\mu s}) \propto f_0^Y \left(\hat{Z}_{l \rightarrow \mu s} + F_{l\mu s} u_{\mu s}, \bar{Z}_{l \rightarrow \mu s} + H_{l\mu s} u_{\mu s} + G_{l\mu s} u_{\mu s}^2 \right). \quad (30)$$

Here, we use the simplified notation $f_i^Y \equiv f_i^{p_Y|Z}$. In appendix A, we show how by making a Taylor expansion of this equation, we can express the message (23) as

$$m_{\mu s \rightarrow l}(u_{\mu s}) \propto p(u_{\mu s}) \mathcal{N}(u_{\mu s}; \hat{U}_{\mu s \rightarrow l}, \bar{U}_{\mu s \rightarrow l}), \quad (31)$$

with

$$\bar{U}_{\mu s \rightarrow l} = - \left(\sum_{l' \neq l} (F_{l' \mu s}^2 + G_{l' \mu s}) \bar{g}_{l' \rightarrow \mu s} + G_{l' \mu s} \hat{g}_{l' \rightarrow \mu s}^2 \right)^{-1}, \quad (32)$$

$$\hat{U}_{\mu s \rightarrow l} = \bar{U}_{\mu s \rightarrow l} \sum_{l' \neq l} F_{l' \mu s} \hat{g}_{l' \rightarrow \mu s}, \quad (33)$$

where

$$\hat{g}_{l' \rightarrow \mu s} = \hat{g}^Y(\hat{Z}_{l' \rightarrow \mu s}, \bar{Z}_{l' \rightarrow \mu s}), \quad \bar{g}_{l' \rightarrow \mu s} = \bar{g}^Y(\hat{Z}_{l' \rightarrow \mu s}, \bar{Z}_{l' \rightarrow \mu s}), \quad (34)$$

and $(\hat{g}^Y(\cdot, \cdot), \bar{g}^Y(\cdot, \cdot))$ are simplified notations for the functions $(\hat{g}^{p_Y|Z}(\cdot, \cdot), \bar{g}^{p_Y|Z}(\cdot, \cdot))$ defined in (20).

This allows us to have a simple expression for the previously introduced mean and variance $\hat{u}_{\mu s \rightarrow l}$ and $\bar{u}_{\mu s \rightarrow l}$ of the message (31). Using the notations (15, 16),

$$\hat{u}_{\mu s \rightarrow l} = \hat{f}^U(\hat{U}_{\mu s \rightarrow l}, \bar{U}_{\mu s \rightarrow l}), \quad \bar{u}_{\mu s \rightarrow l} = \bar{f}^U(\hat{U}_{\mu s \rightarrow l}, \bar{U}_{\mu s \rightarrow l}), \quad (35)$$

where as before, we introduce the simplifying notation $f^U \equiv f^{p_U}$. As noted previously, the exact same thing can be done for the messages (n, \tilde{n}) . The result is an iterative set of equations on a set of means and variances

$$\left(\hat{Z}_{\cdot \rightarrow \cdot}^t, \bar{Z}_{\cdot \rightarrow \cdot}^t, \hat{g}_{\cdot \rightarrow \cdot}^t, \bar{g}_{\cdot \rightarrow \cdot}^t, \hat{U}_{\cdot \rightarrow \cdot}^t, \bar{U}_{\cdot \rightarrow \cdot}^t, \hat{u}_{\cdot \rightarrow \cdot}^t, \bar{u}_{\cdot \rightarrow \cdot}^t, \hat{V}_{\cdot \rightarrow \cdot}^t, \bar{V}_{\cdot \rightarrow \cdot}^t, \hat{v}_{\cdot \rightarrow \cdot}^t, \bar{v}_{\cdot \rightarrow \cdot}^t \right) \quad (36)$$

that constitute the message-passing algorithm.

This algorithm can be further simplified using the so-called Thouless-Andersen-Palmer (TAP) approximation introduced in the study of spin glasses [26]. We refer the reader to other works in which these simplifications are treated in details [10, 19] and only give the resulting algorithm 1, in which only local quantities and no messages are updated. This algorithm is a special case of the ‘‘P-BiG-AMP’’ algorithm, introduced in [19].

As its counterparts for generalized linear models (GAMP [22]) or matrix factorization [10, 21], algorithm 1 needs some adaptations that improve its convergence. One very simple damping scheme that allows to improve convergence (though not guaranteeing it) consists in damping a single variable:

$$\hat{\mathbf{U}}_{t+1} \leftarrow \beta \hat{\mathbf{U}}_{t+1} + (1 - \beta) \hat{\mathbf{U}}_t, \quad (37)$$

with $\beta = 0.3$, applied right after the calculation of $\hat{\mathbf{U}}_{t+1}$. A more involved and better performing, adaptive damping strategy is presented in [27]. Notice that we defined the operators \mathcal{A}_U and \mathcal{A}_V used in algorithm 1 as linear applications $\mathcal{A}_U : \mathbb{R}^P \rightarrow \mathbb{R}^{L \times M}$ and $\mathcal{A}_V : \mathbb{R}^M \rightarrow \mathbb{R}^{L \times P}$ in (10,11): In the algorithm, we apply them row-wise on the matrices they act on.

Algorithm 1 P-BiG-AMP for matrix compressed sensing

Initialization:

Initialize the means $(\hat{\mathbf{u}}_0, \hat{\mathbf{v}}_0, \hat{\mathbf{g}}_0)$ and the variances $(\bar{\mathbf{u}}_0, \bar{\mathbf{v}}_0)$ at random according to the distributions p_U^0 and p_V^0 .

Main loop: while $t < t_{\max}$, calculate following quantities:

$$\begin{aligned}
 \bar{\mathbf{X}}_{t+1} &= \bar{\mathbf{u}}_t \bar{\mathbf{v}}_t^\top + \bar{\mathbf{u}}_t (\hat{\mathbf{v}}_t^2)^\top + \hat{\mathbf{u}}_t^2 \bar{\mathbf{v}}_t^\top \\
 \hat{\mathbf{X}}_{t+1} &= \hat{\mathbf{u}}_t \hat{\mathbf{v}}_t^\top \\
 \bar{\mathbf{Z}}_{t+1} &= \mathcal{A}^2(\bar{\mathbf{X}}_{t+1}) \\
 \hat{\mathbf{Z}}_{t+1} &= \mathcal{A}(\hat{\mathbf{X}}_{t+1}) - \hat{\mathbf{g}}_t \odot \left(\bar{\mathbf{u}}_t (\mathcal{A}_U(\hat{\mathbf{v}}_t) \odot \mathcal{A}_U(\hat{\mathbf{v}}_{t-1}))^\top + (\mathcal{A}_V(\hat{\mathbf{u}}_t) \odot \mathcal{A}_V(\hat{\mathbf{u}}_{t-1})) \bar{\mathbf{v}}_t^\top \right) \\
 \bar{\mathbf{g}}_{t+1} &= \bar{g}^Y(\hat{\mathbf{Z}}_{t+1}, \bar{\mathbf{Z}}_{t+1}) \\
 \hat{\mathbf{g}}_{t+1} &= \hat{g}^Y(\hat{\mathbf{Z}}_{t+1}, \bar{\mathbf{Z}}_{t+1}) \\
 \bar{\mathbf{U}}_{t+1} &= - \left([\mathcal{A}_U(\hat{\mathbf{v}}_t)^2 + \mathcal{A}_U^2(\bar{\mathbf{v}}_t)] \bar{\mathbf{g}}_{t+1} + \mathcal{A}_U^2(\bar{\mathbf{v}}_t) \hat{\mathbf{g}}_{t+1}^2 \right)^{-1} \\
 \hat{\mathbf{U}}_{t+1} &= \bar{\mathbf{U}}_{t+1} \odot (\mathcal{A}_U(\hat{\mathbf{v}}_t) \hat{\mathbf{g}}_{t+1} - \hat{\mathbf{u}}_t \odot \mathcal{A}_U(\hat{\mathbf{v}}_t)^2 \bar{\mathbf{g}}_{t+1} - \hat{\mathbf{u}}_{t-1} \odot \mathcal{A}_U^2(\bar{\mathbf{v}}_{t-1}) \hat{\mathbf{g}}_{t+1} \odot \hat{\mathbf{g}}_t) \\
 \bar{\mathbf{u}}_{t+1} &= \bar{f}^U(\hat{\mathbf{U}}_{t+1}, \bar{\mathbf{U}}_{t+1}) \\
 \hat{\mathbf{u}}_{t+1} &= \hat{f}^U(\hat{\mathbf{U}}_{t+1}, \bar{\mathbf{U}}_{t+1}) \\
 \bar{\mathbf{V}}_{t+1} &= - \left([\mathcal{A}_V(\hat{\mathbf{u}}_t)^2 + \mathcal{A}_V^2(\bar{\mathbf{u}}_t)] \bar{\mathbf{g}}_{t+1} + \mathcal{A}_V^2(\bar{\mathbf{u}}_t) \hat{\mathbf{g}}_{t+1}^2 \right)^{-1} \\
 \hat{\mathbf{V}}_{t+1} &= \bar{\mathbf{V}}_{t+1} \odot (\mathcal{A}_V(\hat{\mathbf{u}}_t) \hat{\mathbf{g}}_{t+1} - \hat{\mathbf{v}}_t \odot \mathcal{A}_V(\hat{\mathbf{u}}_t)^2 \bar{\mathbf{g}}_{t+1} - \hat{\mathbf{v}}_{t-1} \odot \mathcal{A}_V^2(\bar{\mathbf{u}}_{t-1}) \hat{\mathbf{g}}_{t+1} \odot \hat{\mathbf{g}}_t) \\
 \bar{\mathbf{v}}_{t+1} &= \bar{f}^V(\hat{\mathbf{V}}_{t+1}, \bar{\mathbf{V}}_{t+1}) \\
 \hat{\mathbf{v}}_{t+1} &= \hat{f}^V(\hat{\mathbf{V}}_{t+1}, \bar{\mathbf{V}}_{t+1})
 \end{aligned}$$

Result : $(\hat{\mathbf{U}}, \hat{\mathbf{V}}, \hat{\mathbf{X}}, \hat{\mathbf{Z}})$ are the estimates for $(\mathbf{U}, \mathbf{V}, \mathbf{X}, \mathbf{Z})$ and $(\bar{\mathbf{U}}, \bar{\mathbf{V}}, \bar{\mathbf{X}}, \bar{\mathbf{Z}})$ are variances of these estimates.

III. ASYMPTOTIC ANALYSIS

The problem of low-rank matrix compressed sensing can be analyzed with statistical physics methods in the thermodynamic limit, i.e. when the dimensions of the signals M and P and of the measurements L go to infinity. R can remain finite or go to infinity as well. On the other hand, the ratios defined in (7) have to be fixed and finite. As in related inference problems, the analysis is done with the replica method. The resulting state evolution equations describe the behavior of the corresponding message-passing algorithm. In this section, we will focus on the derivation of the replica analysis that results in a simple set of state evolution equations. The analysis is very similar to the one of related inference problems [8, 10, 11, 18].

A. Replica analysis: free entropy

Treating an inference problem as a statistical physics problem reduces to writing an energy function corresponding to the problem and studying the free energy of the system. We are thus interested in calculating a partition function. Here, the relevant partition function is the normalization constant of the probability distribution (21):

$$\mathcal{Z}(\mathbf{Y}, \mathbf{A}) = \int d\mathbf{U} p_U(\mathbf{U}) \int d\mathbf{V} p_V(\mathbf{V}) \int d\mathbf{z} p_{Y|Z}(\mathbf{Y}|\mathbf{z}) \delta[\mathbf{z} - \mathcal{A}(\mathbf{U}\mathbf{V}^\top)]. \quad (38)$$

The free entropy $\log \mathcal{Z}(\mathbf{Y}, \mathbf{A})$ of a given instance can be calculated from the marginals calculated by the belief propagation equations.

However, one can also be interested in the average free entropy of this problem. In order to do this, one needs to average $\log \mathcal{Z}(\mathbf{Y}, \mathbf{A})$ over all possible realizations of \mathbf{A} and \mathbf{Y} , for which we use the replica method [16, 18]. It uses the identity

$$\langle \log \mathcal{Z} \rangle = \lim_{n \rightarrow 0} \frac{\partial}{\partial n} \langle \mathcal{Z}^n \rangle \quad (39)$$

where $\langle \cdot \rangle$ denotes the average over \mathbf{A} and \mathbf{Y} , and relies on the fact that an expression for \mathcal{Z}^n can be found for integer n . This expression is then used for calculating the $n \rightarrow 0$ limit in (39). Though not rigorous, this method has proven to give correct results in a wide range of problems [16, 18].

Let us therefore start by calculating

$$\mathcal{Z}(\mathbf{Y}, \mathbf{A})^n = \int \prod_{a=1}^n \{d\mathbf{U}^a p_U(\mathbf{U}^a) d\mathbf{V}^a p_V(\mathbf{V}^a) d\mathbf{z}^a p_{Y|Z}(\mathbf{Y}|\mathbf{z}^a) \delta[\mathbf{z}^a - \mathcal{A}(\mathbf{U}^a(\mathbf{V}^a)^\top)]\} \quad (40)$$

and its average with respect to the realizations of \mathbf{Y} , generated by \mathbf{U}^0 , \mathbf{V}^0 and \mathcal{A} :

$$\begin{aligned} \langle \mathcal{Z}^n \rangle = & \int d\mathbf{U}^0 p_U^0(\mathbf{U}^0) d\mathbf{V}^0 p_V^0(\mathbf{V}^0) d\mathbf{A} p_A^0(\mathbf{A}) d\mathbf{Y} \\ & d\mathbf{z}^0 p_{Y|Z}(\mathbf{Y}|\mathbf{z}^0) \delta[\mathbf{z}^0 - \mathcal{A}(\mathbf{U}^0(\mathbf{V}^0)^\top)] \mathcal{Z}(\mathbf{Y}, \mathbf{A})^n. \end{aligned} \quad (41)$$

The indices a represent so-called replicas of the system and are initially independent from each other. Carrying on the calculation requires to couple them. To be more precise, each variable $z_l^a = [\mathcal{A}(\mathbf{U}^a(\mathbf{V}^a)^\top)]_l$ is the sum of a large number of independent random variables and can therefore be approximated as a Gaussian random variable. This was done in section II A already and allows again to considerably reduce the number of integrals caused by the averaging over \mathbf{A} . However, z_l^a and z_l^b are not independent, as they are produced with the same operator \mathcal{A} . We show in appendix B that $\mathbf{z}_l \equiv (z_l^0 \dots z_l^n)$ is a multivariate random Gaussian variable with mean 0 and covariance matrix $\mathbf{Q}_z \equiv \mathbf{Q}_u \odot \mathbf{Q}_v$, where the elements of the matrices \mathbf{Q}_u and \mathbf{Q}_v are given by:

$$Q_u^{ab} \equiv \frac{1}{M} \sum_{\mu} u_{\mu}^a u_{\mu}^b, \quad Q_v^{ab} \equiv \frac{1}{P} \sum_p v_p^a v_p^b. \quad (42)$$

As in (41), these quantities can be anything, we have to integrate over them, such that

$$\begin{aligned} \langle \mathcal{Z}^n \rangle = & \int d\mathbf{Q}_u \left[\int \prod_a d\mathbf{U}^a p_U^a(\mathbf{U}^a) \prod_{a \leq b} \delta \left(M Q_u^{ab} - \sum_{\mu} u_{\mu}^a u_{\mu}^b \right) \right] \\ & \int d\mathbf{Q}_v \left[\int \prod_a d\mathbf{V}^a p_V^a(\mathbf{V}^a) \prod_{a \leq b} \delta \left(P Q_v^{ab} - \sum_p v_p^a v_p^b \right) \right] \\ & \prod_{l=1}^L \left[\int d\mathbf{z}_l \mathcal{N}(\mathbf{z}_l; 0, \mathbf{Q}_z) \int d\mathbf{y}_l p_{Y|Z}^0(\mathbf{y}_l|\mathbf{z}_l^0) \prod_{a=1}^n p_{Y|Z}(\mathbf{y}_l|\mathbf{z}_l^a) \right]. \end{aligned} \quad (43)$$

Here, we use the convention that $p_U^a = p_U$ if $a \neq 0$. We now see that the different replicas are coupled via \mathbf{Q}_u and \mathbf{Q}_v in the first two lines. As we did with \mathbf{z}_l , we now introduce the vector $\mathbf{u}_{ps} = (u_{ps}^0 \dots u_{ps}^n)$ (similarly for \mathbf{v}_{ps}) and we use the integral representation of the δ function, introducing the conjugate variables $\hat{\mathbf{Q}}_u$ and $\hat{\mathbf{Q}}_v$ (details in appendix B), which leads to

$$\begin{aligned} \langle \mathcal{Z}^n \rangle = & \int d\mathbf{Q}_u d\hat{\mathbf{Q}}_u e^{-\frac{MR}{2} \text{Tr}(\mathbf{Q}_u \hat{\mathbf{Q}}_u)} \left[\prod_{\mu s} d\mathbf{u}_{\mu s} p_u(\mathbf{u}_{\mu s}) e^{\frac{1}{2} \mathbf{u}_{\mu s}^\top \hat{\mathbf{Q}}_u \mathbf{u}_{\mu s}} \right] \\ & \int d\mathbf{Q}_v d\hat{\mathbf{Q}}_v e^{-\frac{PR}{2} \text{Tr}(\mathbf{Q}_v \hat{\mathbf{Q}}_v)} \left[\prod_{ps} d\mathbf{v}_{ps} p_v(\mathbf{v}_{ps}) e^{\frac{1}{2} \mathbf{v}_{ps}^\top \hat{\mathbf{Q}}_v \mathbf{v}_{ps}} \right] \\ & \prod_{l=1}^L \left[\int d\mathbf{z}_l \mathcal{N}(\mathbf{z}_l; 0, \mathbf{Q}_z) \int d\mathbf{y}_l p_{Y|Z}^0(\mathbf{y}_l|\mathbf{z}_l^0) \prod_{a=1}^n p_{Y|Z}(\mathbf{y}_l|\mathbf{z}_l^a) \right]. \end{aligned} \quad (44)$$

Finally, we assume the distributions of $u_{\mu s}$'s, v_{ps} 's and y_l 's are the same for every coordinate. Using the notations

$$p_u(\mathbf{u}) = p_U^0(u^0) \prod_{a>0} p_U(u^a), \quad p_v(\mathbf{v}) = p_V^0(v^0) \prod_{a>0} p_V(v^a), \quad p_{y|\mathbf{z}}(y|\mathbf{z}) = p_{Y|Z}^0(y|z^0) \prod_{a>0} p_{Y|Z}(y|z^a), \quad (45)$$

this leads to:

$$\begin{aligned} \langle \mathcal{Z}^n \rangle &= \int d\mathbf{Q}_u d\hat{\mathbf{Q}}_u e^{-\frac{MR}{2}\text{Tr}(\mathbf{Q}_u \hat{\mathbf{Q}}_u)} \left[\int d\mathbf{u} p_u(\mathbf{u}) e^{\frac{1}{2}\mathbf{u}^\top \hat{\mathbf{Q}}_u \mathbf{u}} \right]^{RM} \\ &\quad \int d\mathbf{Q}_v d\hat{\mathbf{Q}}_v e^{-\frac{PR}{2}\text{Tr}(\mathbf{Q}_v \hat{\mathbf{Q}}_v)} \left[\int d\mathbf{v} p_v(\mathbf{v}) e^{\frac{1}{2}\mathbf{v}^\top \hat{\mathbf{Q}}_v \mathbf{v}} \right]^{RP} \\ &\quad \left[\int d\mathbf{z} \mathcal{N}(\mathbf{z}; 0, \mathbf{Q}_z) \int dy p_{y|\mathbf{z}}(y|\mathbf{z}) \right]^L. \end{aligned} \quad (46)$$

In the “thermodynamic” limit, we take M , P and L going to infinity with constant ratios. This motivates us to rewrite the last equation under the form

$$\langle \mathcal{Z}^n \rangle = \int d\mathbf{Q}_u d\hat{\mathbf{Q}}_u d\mathbf{Q}_v d\hat{\mathbf{Q}}_v e^{-MR[S_n(\mathbf{Q}_u, \hat{\mathbf{Q}}_u, \mathbf{Q}_v, \hat{\mathbf{Q}}_v)]} \quad (47)$$

and to use the saddle point method, according to which

$$\log(\langle \mathcal{Z}^n \rangle) = -MR \min_{\mathbf{Q}_u, \hat{\mathbf{Q}}_u, \mathbf{Q}_v, \hat{\mathbf{Q}}_v} S_n(\mathbf{Q}_u, \hat{\mathbf{Q}}_u, \mathbf{Q}_v, \hat{\mathbf{Q}}_v). \quad (48)$$

We are therefore left with a minimization problem over the space of the matrices \mathbf{Q}_u , $\hat{\mathbf{Q}}_u$, \mathbf{Q}_v and $\hat{\mathbf{Q}}_v$, representing $2(n+1)(n+2)$ parameters (as the matrices are symmetric).

B. Replica symmetric assumption

The idea of the replica symmetric assumption is that the n replicas introduced in (40) are all equivalent, as they are purely a mathematical manipulation. Based on this, we make the assumption that a sensible matrix \mathbf{Q}_u does not make any distinction between the n introduced replicas. We therefore parametrize \mathbf{Q}_u and $\hat{\mathbf{Q}}_u$ in the following way:

$$\mathbf{Q}_u = \begin{pmatrix} Q_u^0 & m_u & \cdots & m_u \\ m_u & Q_u & \cdots & q_u \\ \vdots & \vdots & \ddots & \vdots \\ m_u & q_u & \cdots & Q_u \end{pmatrix} \quad \hat{\mathbf{Q}}_u = \begin{pmatrix} \hat{Q}_u^0 & \hat{m}_u & \cdots & \hat{m}_u \\ \hat{m}_u & \hat{Q}_u & \cdots & \hat{q}_u \\ \vdots & \vdots & \ddots & \vdots \\ \hat{m}_u & \hat{q}_u & \cdots & \hat{Q}_u \end{pmatrix} \quad (49)$$

and similarly for \mathbf{Q}_v , allowing to be left with 16 instead of $2(n+1)(n+2)$ parameters over which to perform the extremization (48). Furthermore, Q_u^0 and Q_v^0 are in fact known, as they are the second moments of the priors p_U^0 and p_V^0 , and therefore we set

$$\hat{Q}_u^0 = 0, \quad \hat{Q}_v^0 = 0, \quad (50)$$

and thus the extremization is only over 12 variables: $(m_u, \hat{m}_u, q_u, \hat{q}_u, Q_u, \hat{Q}_u)$ and $(m_v, \hat{m}_v, q_v, \hat{q}_v, Q_v, \hat{Q}_v)$.

Let us now look in more details at the function S_n to extremize:

$$\begin{aligned} S_n(\mathbf{Q}_u, \mathbf{Q}_v, \hat{\mathbf{Q}}_u, \hat{\mathbf{Q}}_v) &\equiv \left[\frac{1}{2} \text{Tr} \mathbf{Q}_u \hat{\mathbf{Q}}_u - \log \left(\int d\mathbf{u} p_u(\mathbf{u}) e^{\frac{1}{2}\mathbf{u}^\top \hat{\mathbf{Q}}_u \mathbf{u}} \right) \right] \\ &\quad + \frac{M}{P} \left[\frac{1}{2} \text{Tr} \mathbf{Q}_v \hat{\mathbf{Q}}_v - \log \left(\int d\mathbf{v} p_v(\mathbf{v}) e^{\frac{1}{2}\mathbf{v}^\top \hat{\mathbf{Q}}_v \mathbf{v}} \right) \right] \\ &\quad - \frac{L}{RP} \log \left(\int d\mathbf{z} \mathcal{N}(\mathbf{z}; 0, \mathbf{Q}_z) \int dy p_{y|\mathbf{z}}(y|\mathbf{z}) \right). \end{aligned} \quad (51)$$

Thanks to the parametrization (49), the different terms have simple expressions. The traces can simply be written as

$$\text{Tr} \mathbf{Q}_u \hat{\mathbf{Q}}_u = 2nm\hat{m}_u + nQ_u\hat{Q}_u + n(n-1)q_u\hat{q}_u, \quad (52)$$

while we can use that

$$\mathbf{u}^\top \hat{\mathbf{Q}}_u \mathbf{u} = \hat{Q}_u^0 (u^0)^2 + (\hat{Q}_u - \hat{q}_u) \sum_{a>0} (u^a)^2 + \hat{q}_u \left(\sum_{a>0} u^a \right)^2 + 2\hat{m}_u u^0 \sum_{a>0} u^a \quad (53)$$

and the Gaussian transformation $e^{\lambda\alpha^2} = \int \mathrm{D}x e^{\alpha\sqrt{2\lambda}x}$ in order to rewrite the integral $\int \mathrm{d}\mathbf{u} P_u(\mathbf{u}) e^{\frac{1}{2}\mathbf{u}^\top \hat{\mathbf{Q}}_u \mathbf{u}}$ as

$$\mathcal{I}_U^n = \int \mathrm{D}t \int \mathrm{d}u^0 p_U^0(u^0) \left[\int \mathrm{d}u p_U(u) e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0)u} \right]^n. \quad (54)$$

The third line in (51) can be simplified as well. The first step consists in writing the coupled Gaussian random variables $z^0 \dots z^n$ as a function of n independent, standard Gaussian random variables x^a ($a \in [1, n]$) and one additional standard random variable t that couples them all:

$$z^0 = \sqrt{Q_z^0 - \frac{m_z^2}{q_z}} x^0 + \frac{m_z}{\sqrt{q_z}} t, \quad z^a = \sqrt{Q_z - q_z} x^a + \sqrt{q_z} t. \quad (55)$$

Making the change of variables in the integral, we obtain the following expression for $\int \mathrm{d}\mathbf{z} \mathcal{N}(\mathbf{z}; 0, \mathbf{Q}_z) \int \mathrm{d}y P_{y|z}(y|\mathbf{z})$:

$$\mathcal{I}_Z^n = \int \mathrm{d}y \int \mathrm{D}t \left[\int \mathrm{D}x^0 p_{Y|Z}^0(y|\sqrt{Q_z^0 - \frac{m_z^2}{q_z}} x^0 + \frac{m_z}{\sqrt{q_z}} t) \right] \left[\int \mathrm{D}x p_{Y|Z}(y|\sqrt{Q_z - q_z} x + \sqrt{q_z} t) \right]^n. \quad (56)$$

Looking back at the replica trick (39), we have to study the quantity $\lim_{n \rightarrow 0} \frac{\partial}{\partial n} S_n$ and therefore the quantities

$$\mathcal{I}_U(\hat{\mathbf{Q}}) = \lim_{n \rightarrow 0} \frac{\partial}{\partial n} \log \mathcal{I}_U^n = \int \mathrm{D}t \left[\int \mathrm{d}u^0 p_U^0(u^0) \log \left[\int \mathrm{d}u p_U(u) e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0)u} \right] \right], \quad (57)$$

as well as its equivalent \mathcal{I}_V (obtained by replacing all us by vs in (57)) and

$$\mathcal{I}_Z(\mathbf{Q}) = \lim_{n \rightarrow 0} \frac{\partial}{\partial n} \log \mathcal{I}_Z^n = \int \mathrm{d}y \int \mathrm{D}t f_0^{Y,0} \left(\frac{m}{\sqrt{q}} t, Q^0 - \frac{m^2}{q} \right) \log \left(f_0^Y(\sqrt{q}t, Q - q) \right), \quad (58)$$

where $f_i^{Y,0} \equiv f_i^{p_{Y|Z}^0}$. In the end, we obtain the free entropy ϕ as an extremum

$$\begin{aligned} \phi = -\text{extr} \left\{ \left(m_u \hat{m}_u + \frac{1}{2} Q_u \hat{Q}_u - \frac{1}{2} q_u \hat{q}_u - \mathcal{I}_U(\hat{\mathbf{Q}}_u) \right) \right. \\ \left. + \frac{M}{P} \left(m_v \hat{m}_v + \frac{1}{2} Q_v \hat{Q}_v - \frac{1}{2} q_v \hat{q}_v - \mathcal{I}_V(\hat{\mathbf{Q}}_v) \right) - \frac{L}{RP} \mathcal{I}_Z(\mathbf{Q}_u \odot \mathbf{Q}_v) \right\} \end{aligned} \quad (59)$$

over a set of 12 variables. Note that the shift from a minimum in (48) to an extremum in the equation above is a consequence to the hazardous $n \rightarrow 0$ limit in the replica method.

1. Equivalence to generalized matrix factorization

It is interesting to notice that if $L = MP$ and $R = O(M)$, this free entropy is the same as in generalized matrix factorization [10]. This is not an entirely obvious fact, as the two problems are different and that they are identical only if \mathcal{A} is the identity: in generalized matrix factorization, $\mathbf{Z} = \mathbf{X}$.

In order to perform the theoretical analysis of generalized matrix *factorization* as in [10], it is important to take the limit $R \rightarrow \infty$. In fact, it is this limit that ensures that each entry of \mathbf{Z} is the sum of a large number of random variables, which allows to consider that it has a Gaussian distribution. This is a condition both in the derivation of the message-passing algorithm and in the replica analysis. For that reason, generalized matrix factorization with finite R leads to different algorithms and theoretical bounds [14, 15]. However, in matrix compressed sensing, the mixing of coefficients with \mathcal{A} ensures that even if $R = 1$, each element of \mathbf{Z} can be considered to have a Gaussian distribution. Thanks to this, both the algorithm and the analysis are the same, independently of R . Note that it would be natural to write the free entropy (59) with no explicit R -dependence by introducing a global measurement ratio $\alpha \equiv \frac{L}{R(M+P)}$.

Let us examine the case in which $L = MP$ and $R = O(M)$ and the two problems are strictly equivalent. What differentiates the generalized matrix compressed sensing from the generalized matrix factorization case is that \mathcal{A} is not the identity. However, as \mathcal{A} 's coefficients are Gaussian i.i.d., it is with high probability a bijection when $L = MP$, and in this sense the mixing step does not introduce any further difficulty into the problem compared to matrix factorization. If $L > MP$, matrix compressed sensing is not “compressive” and therefore easier than the corresponding matrix factorization problem, because more measurements are available. If $L < MP$, matrix compressed sensing is “compressive”.

C. State evolution equations

In the previous section, we have derived an expression of the free entropy as an extremum of an action function over a set of parameters. In this section, we find self-consistent equations that hold at the values of these parameters extremizing the action. Furthermore, these self-consistent equations can be iterated in order to numerically obtain the extrema of the action.

In order to find the extremum in (59), we simply set all the partial derivatives of ϕ to 0. The difficult part is finding expressions for the derivatives of the integrals \mathcal{I}_U , \mathcal{I}_V and \mathcal{I}_Z , which we detail here. First we do the calculation for \mathcal{I}_U .

$$\begin{aligned}\frac{\partial}{\partial \hat{Q}_u} \mathcal{I}_U(\hat{\mathbf{Q}}_u) &= \int Dt \int du^0 p_U^0(u^0) \frac{\int du p_U(u) u^2 e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0) u}}{\int du p_U(u) e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0) u}}, \\ \frac{\partial}{\partial \hat{q}_u} \mathcal{I}_U(\hat{\mathbf{Q}}_u) &= \int Dt \int du^0 p_U^0(u^0) \frac{\int du p_U(u) \left(-\frac{u^2}{2} + \frac{tu}{2\sqrt{\hat{q}_u}}\right) e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0) u}}{\int du p_U(u) e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0) u}}, \\ \frac{\partial}{\partial \hat{m}_u} \mathcal{I}_U(\hat{\mathbf{Q}}_u) &= \int Dt \int du^0 u^0 p_U^0(u^0) \frac{\int du p_U(u) u e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0) u}}{\int du p_U(u) e^{\frac{\hat{Q}_u - \hat{q}_u}{2} u^2 + (t\sqrt{\hat{q}_u} + \hat{m}_u u^0) u}}.\end{aligned}\quad (60)$$

If we inject these expressions into the extremization equations of ϕ with respect to \hat{Q}_u , \hat{q}_u , \hat{m}_u and use the update functions defined in (14)-(16), we obtain

$$m_u = \int Dt \int du^0 u^0 p_U^0(u^0) \hat{f}^U \left(\frac{\sqrt{\hat{q}_u} t + \hat{m}_u u^0}{\hat{q}_u - \hat{Q}_u}, \frac{1}{\hat{q}_u - \hat{Q}_u} \right), \quad (61)$$

$$Q_u - q_u = \frac{1}{\sqrt{\hat{q}_u}} \int Dt t \int du^0 p_U^0(u^0) \hat{f}^U \left(\frac{\sqrt{\hat{q}_u} t + \hat{m}_u u^0}{\hat{q}_u - \hat{Q}_u}, \frac{1}{\hat{q}_u - \hat{Q}_u} \right), \quad (62)$$

$$Q_u = \int Dt \int du^0 p_U^0(u^0) \left[\hat{f}^U \left(\frac{\sqrt{\hat{q}_u} t + \hat{m}_u u^0}{\hat{q}_u - \hat{Q}_u}, \frac{1}{\hat{q}_u - \hat{Q}_u} \right) + \left(\hat{f}^U \left(\frac{\sqrt{\hat{q}_u} t + \hat{m}_u u^0}{\hat{q}_u - \hat{Q}_u}, \frac{1}{\hat{q}_u - \hat{Q}_u} \right) \right)^2 \right]. \quad (63)$$

These equations can be further simplified by using the transformation $t \leftarrow t + \frac{\hat{m}}{\sqrt{\hat{q}}} u^0$ and integrating by part eq (62):

$$m_u = \sqrt{\frac{\hat{q}_u}{\hat{m}_u^2}} \int dt f_1^{U,0} \left(\frac{\sqrt{\hat{q}_u} t}{\hat{m}_u}, \frac{\hat{q}_u}{\hat{m}_u^2} \right) \hat{f}^U \left(\frac{\sqrt{\hat{q}_u} t}{\hat{q}_u - \hat{Q}_u}, \frac{1}{\hat{q}_u - \hat{Q}_u} \right), \quad (64)$$

$$Q_u - q_u = \sqrt{\frac{\hat{q}_u}{\hat{m}_u^2}} \int dt f_0^{U,0} \left(\frac{\sqrt{\hat{q}_u} t}{\hat{m}_u}, \frac{\hat{q}_u}{\hat{m}_u^2} \right) \hat{f}^U \left(\frac{\sqrt{\hat{q}_u} t}{\hat{q}_u - \hat{Q}_u}, \frac{1}{\hat{q}_u - \hat{Q}_u} \right), \quad (65)$$

$$q_u = \sqrt{\frac{\hat{q}_u}{\hat{m}_u^2}} \int dt f_0^{U,0} \left(\frac{\sqrt{\hat{q}_u} t}{\hat{m}_u}, \frac{\hat{q}_u}{\hat{m}_u^2} \right) \left[\hat{f}^U \left(\frac{\sqrt{\hat{q}_u} t}{\hat{q}_u - \hat{Q}_u}, \frac{1}{\hat{q}_u - \hat{Q}_u} \right) \right]^2, \quad (66)$$

and the same equations hold replacing u by v .

Let us now come to the derivatives of \mathcal{I}_Z . To calculate them, we use the identity (19), taking $s = q$ or $s = \frac{m^2}{q}$. After an integration by parts, we obtain

$$\frac{\partial}{\partial m} \mathcal{I}_Z(\mathbf{Q}) = \frac{1}{m} \int dy \int Dt \frac{\left[\frac{\partial}{\partial t} f_0^{Y,0} \left(\frac{m}{\sqrt{q}} t, Q^0 - \frac{m^2}{q} \right) \right] \left[\frac{\partial}{\partial t} f_0^Y(\sqrt{qt}, Q - q) \right]}{f_0^Y(\sqrt{qt}, Q - q)}, \quad (67)$$

$$\frac{\partial}{\partial q} \mathcal{I}_Z(\mathbf{Q}) = -\frac{1}{2q} \int dy \int Dt \left[\frac{\frac{\partial}{\partial t} f_0^Y(\sqrt{qt}, Q - q)}{f_0^Y(\sqrt{qt}, Q - q)} \right]^2 f_0^{Y,0} \left(\frac{m}{\sqrt{q}} t, Q^0 - \frac{m^2}{q} \right), \quad (68)$$

$$\frac{\partial}{\partial Q} \mathcal{I}_Z(\mathbf{Q}) = \int dy \int Dt f_0^{Y,0} \left(\frac{m}{\sqrt{q}} t, Q^0 - \frac{m^2}{q} \right) \left[\frac{\frac{\partial}{\partial Q} f_0^Y(\sqrt{qt}, Q - q)}{f_0^Y(\sqrt{qt}, Q - q)} \right]. \quad (69)$$

Injecting these expressions into the extremization equations of ϕ with respect to Q, q, m , we obtain

$$\hat{m} = \frac{1}{m} \int dy \int Dt \frac{\left[\frac{\partial}{\partial t} f_0^{Y,0} \left(\frac{m}{\sqrt{q}} t, Q^0 - \frac{m^2}{q} \right) \right] \left[\frac{\partial}{\partial t} f_0^Y (\sqrt{qt}, Q - q) \right]}{f_0^Y (\sqrt{qt}, Q - q)}, \quad (70)$$

$$\hat{q} = \frac{1}{q} \int dy \int Dt \left[\frac{\frac{\partial}{\partial t} f_0^Y (\sqrt{qt}, Q - q)}{f_0^Y (\sqrt{qt}, Q - q)} \right]^2 f_0^{Y,0} \left(\frac{m}{\sqrt{q}} t, Q^0 - \frac{m^2}{q} \right), \quad (71)$$

$$\hat{Q} = 2 \int dy \int Dt f_0^{Y,0} \left(\frac{m}{\sqrt{q}} t, Q^0 - \frac{m^2}{q} \right) \left[\frac{\frac{\partial}{\partial Q} f_0^Y (\sqrt{qt}, Q - q)}{f_0^Y (\sqrt{qt}, Q - q)} \right], \quad (72)$$

and remembering that $m = m_u m_v, q = q_u q_v, Q = Q_u Q_v$ and the definitions (7):

$$\hat{m}_u = \alpha_U m_v \hat{m}, \quad \hat{q}_u = \alpha_U q_v \hat{q}, \quad \hat{Q}_u = \alpha_U Q_v \hat{Q}, \quad (73)$$

$$\hat{m}_v = \alpha_V m_u \hat{m}, \quad \hat{q}_v = \alpha_V q_u \hat{q}, \quad \hat{Q}_v = \alpha_V Q_u \hat{Q}. \quad (74)$$

The equations (64,65,66) along with their equivalents for v , the equations (70,71,72) and (73,74) constitute a closed set of equations that hold at the extrema of ϕ in equation (59).

When they are iterated, they constitute the so-called state evolution equations. These can also be obtained by the analysis of the BP algorithm and are known to accurately describe the algorithm's behavior when the replica symmetric hypothesis is indeed correct.

As noted before, if $L = MP$, these state evolution equations are identical to the ones in matrix factorization [10]. Therefore, they reduce to the state evolution of GAMP when U is known, which corresponds to fixing $m_u = q_u = Q_u = Q_u^0$ in the equations.

D. Bayes-optimal analysis

Until now, we have not supposed exact knowledge of the true signal distributions and of the true measurement channel. When this is the case, the state evolution equations greatly simplify because of the so-called Nishimori conditions [28]. In our case, these ensure that following equalities hold:

$$Q = Q^0, \quad \hat{Q} = 0, \quad m = q, \quad \hat{m} = \hat{q} \quad (75)$$

both for u and v . Then, we only need to keep track of the variables $(m_u, \hat{m}_u, m_v, \hat{m}_v)$, and the state evolution is obtained by choosing initial values for (m_u^0, m_v^0) and iterating for $i \geq 0$ the equations

$$\hat{m}^{i+1} = \frac{1}{m_u^i m_v^i} \int dy \int Dt \frac{\left[\frac{\partial}{\partial t} f_0^Y (\sqrt{m_u^i m_v^i} t, Q_u^0 Q_v^0 - m_u^i m_v^i) \right]^2}{f_0^Y (\sqrt{m_u^i m_v^i} t, Q_u^0 Q_v^0 - m_u^i m_v^i)}, \quad (76)$$

$$m_u^{i+1} = \frac{1}{\sqrt{\alpha_U m_v^i \hat{m}^{i+1}}} \int dt \frac{\left[f_1^U \left(\frac{t}{\sqrt{\alpha_U m_v^i \hat{m}^{i+1}}}, \frac{1}{\alpha_U m_v^i \hat{m}^{i+1}} \right) \right]^2}{f_0^U \left(\frac{t}{\sqrt{\alpha_U m_v^i \hat{m}^{i+1}}}, \frac{1}{\alpha_U m_v^i \hat{m}^{i+1}} \right)}, \quad (77)$$

$$m_v^{i+1} = \frac{1}{\sqrt{\alpha_V m_u^i \hat{m}^{i+1}}} \int dt \frac{\left[f_1^V \left(\frac{t}{\sqrt{\alpha_V m_u^i \hat{m}^{i+1}}}, \frac{1}{\alpha_V m_u^i \hat{m}^{i+1}} \right) \right]^2}{f_0^V \left(\frac{t}{\sqrt{\alpha_V m_u^i \hat{m}^{i+1}}}, \frac{1}{\alpha_V m_u^i \hat{m}^{i+1}} \right)}, \quad (78)$$

until convergence. From m_u and m_v , one can simply deduce the mean squared errors by the following relations:

$$\text{MSE}_u = Q_u^0 - m_u, \quad \text{MSE}_v = Q_v^0 - m_v, \quad \text{MSE}_x = Q_u^0 Q_v^0 - m_u m_v. \quad (79)$$

The initialization values $(m_u, \hat{m}_u, m_v, \hat{m}_v)$ indicate how close to the solution the algorithm is at initialization. In case of a random initialization of the algorithm, the expected initial overlaps m_u^0 and m_v^0 are of order $1/M$ and $1/P$ respectively, and they should therefore be set to these values (or less) in the state evolution equations.

Note that state evolution run with matching priors without imposing the Nishimori conditions (75) should in principle give the exact same results as the Bayes-optimal state evolution analysis presented above, and thus naturally

follow the so-called “Nishimori line” defined by (75). However, as shown in [2], the Nishimori line can be unstable: In that case, numerical fluctuations around it will be amplified under iterations of state evolution that will thus give a different result than its counterpart with imposed Nishimori conditions. This instability of the Nishimori line seems to be the reason why algorithm 1 as well as others of the same type do not converge without damping of the variables.

IV. CASE STUDY

In this section, we focus on one specific setting for which the state evolution equations are practical to implement. An analysis of their fixed points leads to an understanding of different phases and of the phase transitions between them.

We look at the setting in which both \mathbf{U} and \mathbf{V} follow a Bernoulli-Gauss distribution:

$$p_U(u) = (1 - \rho_u)\delta(u) + \rho_u \mathcal{N}(u; 0, 1), \quad (80)$$

$$p_V(v) = (1 - \rho_v)\delta(v) + \rho_v \mathcal{N}(v; 0, 1), \quad (81)$$

and the measurements are taken through an additive white Gaussian noise (AWGN) channel:

$$\forall l \in [1, L], \quad Y_l = [\mathbf{A}(\mathbf{U}\mathbf{V}^T)]_l + \xi_l, \quad \text{with } \xi_l \sim \mathcal{N}(\xi_l; 0, \Delta). \quad (82)$$

Note that most previous works [7, 13, 23, 29] consider this channel. For the AWGN channel the equation (76) has a simple analytical expression:

$$\hat{m}^{i+1} = \frac{1}{\Delta + \rho_u \rho_v - m_u^i m_v^i}. \quad (83)$$

Further simplifying the setting to the special case $M = P$ and $\rho_u = \rho_v = \rho$, the Bayes optimal state evolution equations (76-78) can be written as one single equation

$$m = \sqrt{\frac{\Delta + \rho^2 - m^2}{\alpha_u m}} \int dt \frac{\left[f_1^U \left(\sqrt{\frac{\Delta + \rho^2 - m^2}{\alpha_u m}} t, \frac{\Delta + \rho^2 - m^2}{\alpha_u m} \right) \right]^2}{f_0 \left(\sqrt{\frac{\Delta + \rho^2 - m^2}{\alpha_u m}} t, \frac{\Delta + \rho^2 - m^2}{\alpha_u m} \right)}, \quad (84)$$

in which the iteration-time indices of m , i (left hand side) and $i - 1$ (right hand side), are left out for better legibility. We can define a global measurement rate

$$\alpha \equiv \frac{L}{2MR} = \frac{\alpha_u}{2}, \quad (85)$$

which is the natural quantity to compare ρ to.

A. Phases and phase transitions

As in compressed sensing or in matrix factorization, the analysis of the free entropy and state evolution equations reveals the existence of different phases in which the difficulty of the problem is different. In our case study, the free entropy ϕ has the following expression:

$$\begin{aligned} \phi(m) = & -m\hat{m} - \frac{\alpha}{4} \log(2\pi(\Delta + \rho^2 - m^2)) \\ & + \frac{2}{\sqrt{\hat{m}}} \int dt f_0^U \left(\frac{t}{\sqrt{\hat{m}}}, \frac{1}{\hat{m}} \right) \left[\frac{t^2}{2} + \log \left(\sqrt{\frac{2\pi}{\hat{m}}} f_0^U \left(\frac{t}{\sqrt{\hat{m}}}, \frac{1}{\hat{m}} \right) \right) \right] \end{aligned} \quad (86)$$

with

$$\hat{m} = \frac{1}{\Delta + \rho^2 - m^2}. \quad (87)$$

The integral can best be numerically evaluated replacing \int by $2 \left(\int_0^{20} + \int_{20}^{20\sqrt{1+\hat{m}}} \right)$, which allows a reliable numerical evaluation for all possible values of \hat{m} .




FIG. 3. Free entropy landscapes for $\rho = 0.5$, $\Delta = 10^{-5}$. Crosses represent local maxima. There are three types of them: either at $\text{nMSE} = 1$ (as for $\alpha = 0.49$), or at $\text{nMSE} \approx \Delta$, or in an intermediary region. In case there are several local maxima (as for $\alpha = 0.68$), the algorithm will perform sub-optimally, getting stuck in the local maximum of highest nMSE instead of converging to the global maximum (“hard but possible” phase).

Figure 3 shows the free entropy landscapes for $\rho = 0.1$ and different values of α . Instead of using m as x -axes, we use the normalized mean squared error

$$\text{nMSE} = 1 - \frac{m}{\rho}, \quad (88)$$

that is a more natural quantity to measure the quality of reconstruction.

We can define three different phases depending on the positions of the free entropy maxima. In the noiseless setting, these are:

1. An “impossible” phase, in which the global maximum of the free entropy is not at $\text{nMSE} = 0$. In that phase, no algorithm can find the correct solution.
2. A “hard but possible” phase, in which the free entropy has its global maximum at $\text{nMSE} = 0$, but also a *local* maximum at non-zero nMSE . In that phase, it is possible to find the correct solution, by correctly sampling from the posterior distribution (21). However, algorithms such as P-BiG-AMP get stuck in the local free entropy maximum instead of finding the global maximum.
3. An “easy” phase, in which the free entropy function has a single maximum at $\text{nMSE} = 0$.

In a noisy setting as in figure 3, the lowest achievable nMSE is of the order of the AWGN variance Δ instead of 0.

1. State evolution fixed points

The state evolution equation (84) can either be iterated or considered as a fixed point equation. Figure 4 shows the fixed points of (84), which are all local extrema of the free entropy ϕ . The iterated state evolution equation converges to one of the local maxima. Since the state evolution for the matrix compressed sensing problem and the dictionary learning problem are the same (provided $L = MP$ and $R = O(M)$) these diagrams and their analysis are equivalent to those presented in previous work on the dictionary learning [10]. Notably [25] presented analogous diagrams depicting the fixed points for the dictionary learning problem.

The plots allow to see more clearly the “impossible”, “hard but possible” and “easy” phases. In the “hard but possible” phase, the state evolution has an unstable fixed point, which corresponds to a local minimum of the free entropy. Three interesting facts can be noticed:

1. In the noiseless setting, the impossible/possible phase transition (the apparition of the low nMSE fixed point) takes place at $\alpha = \rho$. This can be expected because it is the critical α at which the number of available equations is equal to the total number of non-zero components of the unknowns, just as in compressed sensing.




FIG. 4. Fixed points of the state evolution equation (84) for two different sparsities ρ . For values of α for which two stable fixed points exist, the iterated state evolution equation converges to the one of higher nMSE if the initial nMSE is higher than the unstable fixed point, and to the one of lower nMSE if not.

2. The fixed point at $\text{nMSE}=1$ always exists and is stable for $\alpha \in [0, 1/2]$. This is a rather remarkable fact that does not appear in compressed sensing. A consequence of this is the existence of a “hard but possible” phase that even for very small values of ρ extends at least up to $\alpha = 1/2$. This radically differs from the low- ρ regime in compressed sensing, in which the measurement rate α necessary for tractable recovery goes to zero as $\rho \rightarrow 0$.
3. Increasing α starting below $1/2$ and following the high-nMSE branch, two successive phase transitions are encountered. First, the $\text{nMSE} = 1$ fixed point disappears at $\alpha = 1/2$ and turns into an $\text{nMSE} < 1$ fixed point in a second order (i.e. continuous) phase transition. Second, the upper branch disappears and the discontinuity of the nMSE of the fixed point, jumping down to the lower branch, marks a first order phase transition. While these two transitions of different types are clearly visible in Figure 4b, they are too close together in Figure 4a to be distinguished. They are separated nonetheless, the easy/hard (first order) phase transition always takes place at $\alpha > 1/2$.

Figure 5 shows the full phase diagram for the case-study problem, with the easy, hard and impossible phases. The “uninformed” line is obtained by starting the state evolution starting from $\text{nMSE} = 1 - \epsilon$, with an infinitesimally small ϵ , and defines the transition between the “easy” and the “hard” phase. Interestingly, the entire region with $\alpha < 0.5$ is in the hard phase, even at low values of ρ , due to the existence of the stable fixed point at $\text{nMSE} = 1$. In the “hard” phase, inference is possible provided a good estimation of the signal is already known. The effect of such a partial knowledge can be simulated by running the state evolution equation (84) starting with $\text{nMSE} = 0.9$, leading to the “informed” line, for which $\alpha \rightarrow 0$ when $\rho \rightarrow 0$. The position of this line depends strongly on the starting nMSE.

B. Comparison with algorithmic performances

Figures 6 and 7 presents a comparison of the theoretical fixed point analysis performed above with the actual performances of P-BiG-AMP.

For the experiments, rank $R = 1$ was used. In this setting, the only invariance left is a scaling invariance: if (\mathbf{U}, \mathbf{V}) is the true solution, then for every $\gamma \neq 0$, $(\gamma\mathbf{U}, \frac{1}{\gamma}\mathbf{V})$ is a solution as well. The final nMSE returned by the algorithm takes this invariance into account and is the average of the error on \mathbf{U} and the error on \mathbf{V} :

$$\text{nMSE} = \frac{1}{2} (\text{nMSE}_{\mathbf{u}} + \text{nMSE}_{\mathbf{v}}) \quad (89)$$

which will be compared to the results obtained by the theoretical expression (88). For each instance of the problem, the algorithm was allowed up to 20 restarts from different random initializations to reach a nMSE smaller than 10^{-6} , and the lowest of the reached nMSE was kept.




FIG. 5. Phase diagram for the considered case-study obtained from the state evolution, eq. (84). Noise variance is $\Delta = 10^{-12}$ and success is defined by a final $\text{nMSE} < 10^{-10}$. The disappearing of the state evolution fixed point (or equivalently, of a free entropy maximum) with nMSE of order 1 marks the frontier between the “easy” and the “hard” phase (full line). The dashed line marks the easy/hard phase boundary when an “informed” initialization is provided (see text). The possible/impossible frontier represented corresponds to the noiseless case.

The results show that there is a good agreement between the theory and the performance of P-BiG-AMP: most of the nMSE s reached by P-BiG-AMP correspond to a stable fixed point of the state evolution. The agreement with the theory becomes better with increasing system size. For smaller sizes, the experimental points are more spread around the theoretical fixed points. This can be well understood by analyzing the case of fixed points with $\text{nMSE}=1$. The “meaning” of such fixed points is that the algorithm is unable to estimate the true signals better than at random. In the $M \rightarrow \infty$ limit, the nMSE between the true signals and random signals is 1 with probability 1. For finite values of M however, the nMSE between true and random signals follows a distribution on $[0, 1]$ that gets more peaked on 1 as M increases. This explains the narrowing of the spread of experimental points around the fixed points as M increases.

1. Succeeding in the hard phase: importance of the initialization

An interesting consequence of this finite size effect is that for small M , parts of the “hard” phase are quite easy. The reason is that if the random initialization of the algorithm is such that the nMSE is *smaller* than the nMSE of the *unstable* fixed point, the algorithm naturally converges to the low- nMSE solution. Therefore, running the algorithm from a few different initializations can allow to converge to the correct solution even in the “hard” phase, provided that M is small enough and that the unstable fixed point has a high enough nMSE .

Figure 8 shows that this effect is quite important for $\rho = 0.1$, but nearly inexistent for $\rho = 0.6$. The reason for this is the much higher nMSE of the unstable fixed point for $\rho = 0.1$ than for $\rho = 0.6$.

Remember that in P-BiG-AMP, the initial estimates of \mathbf{U} and \mathbf{V} are random. While in some regions of the phase diagram and with small signal sizes, running the algorithm from several of those random initial estimates might be sufficient, in general it would be preferable to have a procedure that systematically produces good initializations. Previous works stress this fact as well and often rely on an initialization from spectral methods [7, 13, 23, 29].

Another difference between figures 8a and 8b is that in the latter, the algorithm fails for a significant fraction of instances inside the “easy” phase, which is not the case in the former. The fact that the fraction of such failed instances decreases with increasing signal size M seems to indicate that this is as well a finite size effect. Unlike the previously examined finite size effect, this one cannot be explained from the state evolution, as it has a unique fixed

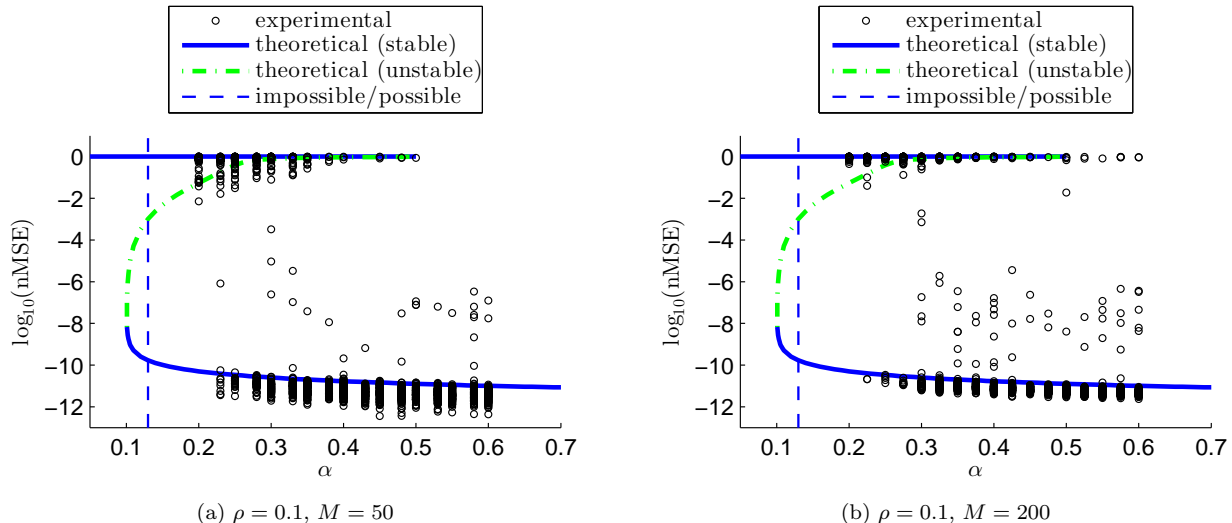


FIG. 6. Comparison of fixed points obtained by the state evolution and values reached by the P-BiG-AMP algorithm. Parameters are $\rho = 0.1$, $\Delta = 10^{-12}$ with (a): $M = 50$, (b): $M = 200$. For each α there are 100 experimental points. The experimental fixed points are relatively close to the fixed points of the state evolution. Note that the spreading around the theoretical line diminishes with growing M . In the thermodynamic limit $M \rightarrow \infty$, all experimental points would be on the fixed point of *highest* nMSE. At finite M , the probability to initialize the algorithm *below* the unstable fixed point allows some instances to converge to the low-nMSE fixed point.

point in the “easy” phase.

V. CONCLUSION

In this paper, we provide an asymptotic analysis of Bayesian low-rank matrix compressed sensing. We employ the replica method of statistical physics to obtain the so-called state evolution equations, whose fixed points allow us to determine if inference is easy, hard or impossible. The state evolution equations describe the behavior of associated message passing algorithm P-BiG-AMP that was derived and studied previously in [19]. This work inscribes in a line of work where approximate message passing was derived and analyzed on related estimation problems such as compressed sensing [6, 22], or matrix factorization [10, 20, 21].

An interesting point concerning the state evolution equations is that they are the same as those for the matrix factorization problem derived in [10]. Related observations were made in [5].

Our analysis, just as the algorithm, is written for a generic separable prior and output channel. We analyze in detail the phase diagram for Gaussian noise on the output and Gauss-Bernoulli prior on both the factors. A striking point in the phase diagram is that the α (eq. (85)) needed for the recovery to be tractable does not go to zero as the factors become very sparse. This is a remarkable difference between the matrix and the linear compressed sensing. We show numerically that there is an excellent agreement between the theoretical analysis and the performances of the P-BiG-AMP algorithm. We observe that for the simulated system sizes, the algorithm performs better than what could be expected from the asymptotic theoretical analysis. However, we explain this as a finite size effect in terms of state evolution fixed points and stress the importance of a good initial estimate in order to perform inference outside of the easy phase. Our analysis quantifies how “good” the initialization needs to be for large systems to allow tractable recovery.

ACKNOWLEDGEMENT

Philip Schniter’s work on this project was supported in part by the National Science Foundation under grant CCF-1527162. Christophe Schülke’s work was supported in part by Université franco-italienne and in part by the ERC under the European Union’s 7th Framework Programme Grant Agreement 307087-SPARCS.

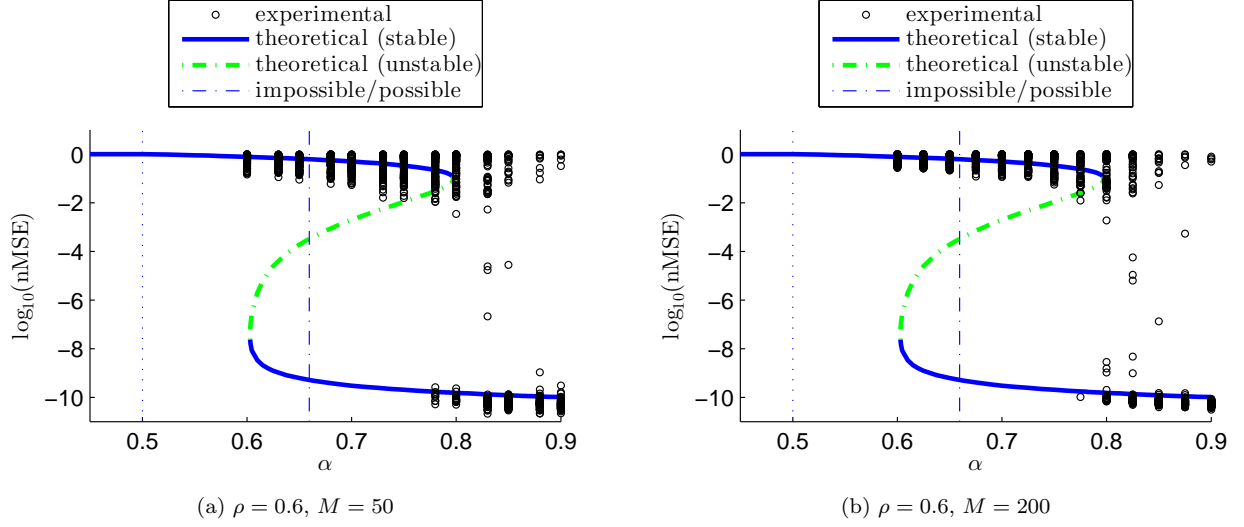


FIG. 7. Comparison of fixed points obtained by the state evolution and values reached by the P-BiG-AMP algorithm. Parameters are $\rho = 0.6$, $\Delta = 3.6 \times 10^{-11}$ with (a): $M = 50$, (b): $M = 200$. For each α there are 100 experimental points. Unlike for the $\rho = 0.1$ case on figure 6, the algorithm fails for an important fraction of instances in the “easy” phase. This phenomenon is not explained by the state evolution analysis and might be a finite size effect. However, as α grows the probability of success goes to 1 (see figure 8b). Unlike for $\rho = 0.1$, the probability of recovery inside the “hard” phase is much smaller, due to the lower nMSE of the unstable fixed point. The thin dotted line marks the position of the second order phase transition, at which the nMSE stops being strictly equal to 1.

Appendix A: Details for the derivation of the message-passing algorithm

Here, we complete the derivation of the message-passing algorithm starting with equation (30):

$$\tilde{m}_{l \rightarrow \mu s}(u_{\mu s}) \propto f_0^Y \left(\hat{Z}_{l \rightarrow \mu s} + F_{l \mu s} u_{\mu s}, \bar{Z}_{l \rightarrow \mu s} + H_{l \mu s} u_{\mu s} + G_{l \mu s} u_{\mu s}^2 \right). \quad (\text{A1})$$

We first make a Taylor expansion of this message at order 2 around $u_{\mu s} = 0$. We drop all indices for this calculation and use simplified notations $f = f_0^Y(\hat{Z}, \bar{Z})$, $\partial_1 = \frac{\partial}{\partial \hat{Z}}$, $\partial_2 = \frac{\partial}{\partial \bar{Z}}$:

$$\begin{aligned} \tilde{m}(u) &\propto f + u(F\partial_1 f + H\partial_2 f) \\ &+ \frac{1}{2}u^2(F^2\partial_1^2 f + H^2\partial_2^2 f + 2FH\partial_1\partial_2 f + 2G\partial_2 f) + o(u^2). \end{aligned} \quad (\text{A2})$$

We can rewrite \tilde{m} as a Gaussian

$$\tilde{m}(u) \propto \mathcal{N}(u; \hat{p}, \bar{p}) + o(u^2) \quad (\text{A3})$$

by identifying the coefficients of the Taylor expansion above with the Taylor expansion of a Gaussian

$$\mathcal{N}(x; \frac{a}{b}, -\frac{1}{b}) \propto 1 - ax + \frac{b+a^2}{2}x^2 + o(x^2). \quad (\text{A4})$$

Note that the form (A3) is only valid around $u = 0$: \tilde{m} is not Gaussian. However this form makes calculations easier. Identification of the coefficients in (A2) and (A4) leads to

$$\begin{aligned} \bar{p} = & - \left[F^2 \left(\frac{\partial_1^2 f}{f} - \left(\frac{\partial_1 f}{f} \right)^2 \right) + 2G \frac{\partial_2 f}{f} \right. \\ & \left. + H^2 \left(\frac{\partial_2^2 f}{f} - \left(\frac{\partial_2 f}{f} \right)^2 \right) + 2FH \left(\frac{\partial_1 \partial_2 f}{f} - \frac{\partial_1 f}{f} \frac{\partial_2 f}{f} \right) \right]^{-1} \end{aligned} \quad (\text{A5})$$

$$\hat{p} = -\bar{p} \left(F \frac{\partial_1 f}{f} + H \frac{\partial_2 f}{f} \right) \quad (\text{A6})$$

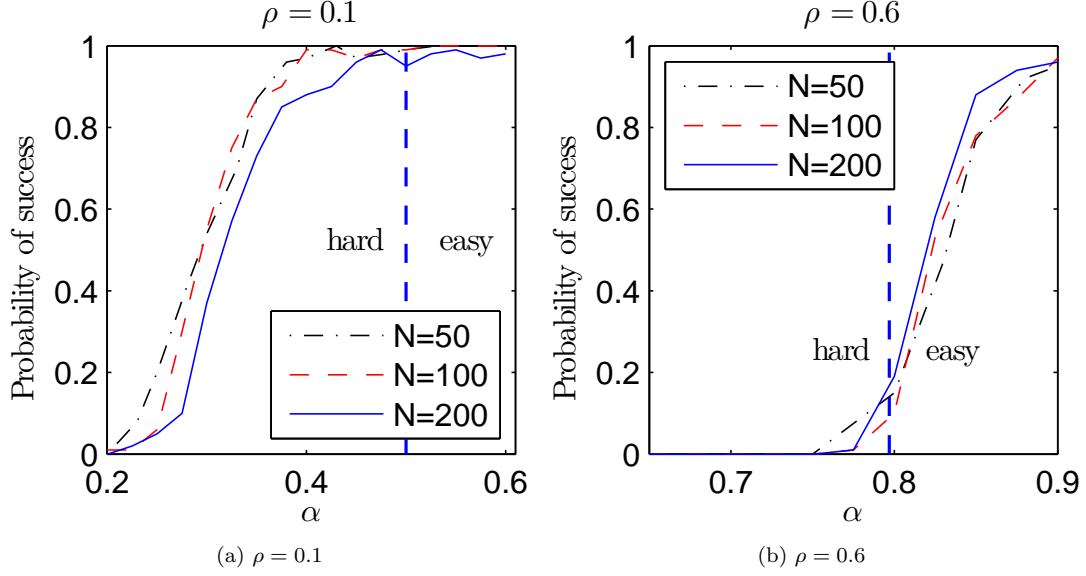


FIG. 8. Empirical probability of success (defined by $\text{nMSE} < 10^{-6}$), for the experiments presented on figures 6 and 7. Due to the finite size, the position of the curves slightly vary for different values of M . Finite size effects allow a fraction of successful instances inside the hard phase for $\rho = 0.1$, but much less for $\rho = 0.6$.

We can now treat the m -messages from eq. (23). The product is easy to handle as it is a product of Gaussians

$$\prod_{l' \neq l} \tilde{m}_{l' \rightarrow \mu s}(u_{\mu s}) \propto \prod_{l' \neq l} \mathcal{N}(u_{\mu s}; \hat{p}_{l' \rightarrow \mu s}, \bar{p}_{l' \rightarrow \mu s}) \propto \mathcal{N}(u_{\mu s}; \hat{U}_{\mu s \rightarrow l}, \bar{U}_{\mu s \rightarrow l}), \quad (\text{A7})$$

which allows us to write

$$\bar{U}_{\mu s \rightarrow l} = \left(\sum_{l' \neq l} \bar{p}_{l' \rightarrow \mu s}^{-1} \right)^{-1} \quad (\text{A8})$$

$$\hat{U}_{\mu s \rightarrow l} = \bar{U}_{\mu s \rightarrow l} \sum_{l' \neq l} \left(\frac{\hat{p}_{l' \rightarrow \mu s}}{\bar{p}_{l' \rightarrow \mu s}} \right) \quad (\text{A9})$$

In the sums above, some of the non-leading order terms stemming from (A5,A6) have a vanishing contribution in the limit where $(M, P, L) \rightarrow \infty$ and will therefore be neglected. The table below analyzes the orders of magnitude and possible signs of all quantities in (A5,A6). In the third and fourth line, we use this to analyze the order of magnitude of a sum of L of those terms, as appears in (A8,A9) and what this leads to when $L \propto RM$, which is the scaling we are interested in.

	F	G	H	F^2	H^2	FH
scales as:	$\frac{1}{\sqrt{RP}}$	$\frac{1}{RP}$	$\frac{1}{R\sqrt{MP}}$	$\frac{1}{RP}$	$\frac{1}{R^2MP}$	$\frac{1}{R^{3/2}P}$
sign:	\pm	$+$	\pm	$+$	$+$	\pm
sum over L	$\frac{\sqrt{L}}{\sqrt{RP}}$	$\frac{L}{RP}$	$\frac{\sqrt{L}}{R\sqrt{MP}}$	$\frac{L}{RP}$	$\frac{L}{R^2MP}$	$\frac{\sqrt{L}}{R^{3/2}P}$
$L \propto RP$	1	1	$\frac{1}{\sqrt{RM}}$	1	$\frac{1}{RM}$	$\frac{1}{R\sqrt{P}}$

This analysis is based on the fact that:

- \mathbf{A} has random i.i.d. elements of mean 0 and variance $1/(RMP)$
- \mathbf{U} , \mathbf{V} and \mathbf{z} have zero-mean elements of order 1, therefore all estimators of type $\hat{\mathbf{u}}$, $\hat{\mathbf{U}}$, etc. are of order 1 as well, either positive or negative
- variances of type $\bar{\mathbf{u}}$, $\bar{\mathbf{U}}$, etc. are positive and of order 1

- all quantities of the type $\frac{\partial_i f}{f}$ are of order 1.

With the help of the table, we can neglect all terms that have a vanishing contribution. Furthermore, using the relations (17,18) and the definition of the g functions (20), it can be shown that

$$\frac{\partial_1 f^Y}{f^Y} = \hat{g}^Y, \quad (\text{A10})$$

$$\frac{\partial_1^2 f^Y}{f^Y} - \left(\frac{\partial_1 f^Y}{f^Y} \right)^2 = \bar{g}^Y, \quad (\text{A11})$$

$$\frac{\partial_2 f^Y}{f^Y} = \frac{1}{2} (\bar{g}^Y - (\hat{g}^Y)^2). \quad (\text{A12})$$

In the end, the resulting expressions for (A8,A9) are given in (32,33).

Appendix B: Details for the replica calculation

a. Covariance matrix of \mathbf{z}_l We treat $z_l^a = [\mathcal{A}(\mathbf{U}^a(\mathbf{V}^a)^\top)]_l$ as a random variable of \mathcal{A} and look at the covariance between two of those variables:

$$\langle z_l^a z_{l'}^b \rangle = \left\langle \left(\sum_{\mu p} A_l^{\mu p} \sum_s u_{\mu s}^a v_{ps}^a \right) \left(\sum_{\mu' p'} A_{l'}^{\mu' p'} \sum_{s'} u_{\mu' s'}^b v_{p's'}^b \right) \right\rangle \quad (\text{B1})$$

$$= \left\langle \sum_{\mu \mu'} \sum_{p p'} A_l^{\mu p} A_{l'}^{\mu' p'} \sum_{s s'} u_{\mu s}^a u_{\mu' s'}^b v_{ps}^a v_{p's'}^b \right\rangle \quad (\text{B2})$$

$$= \sum_{\mu \mu'} \sum_{p p'} \langle A_l^{\mu p} A_{l'}^{\mu' p'} \rangle \sum_{s s'} u_{\mu s}^a u_{\mu' s'}^b v_{ps}^a v_{p's'}^b \quad (\text{B3})$$

As the elements of \mathbf{A} are i.i.d. with zero mean and variance $1/(RMP)$, we have $\langle A_l^{\mu p} A_{l'}^{\mu' p'} \rangle = \delta_{l,l'} \delta_{\mu,\mu'} \delta_{p,p'} \frac{1}{RMP}$ and thus

$$\langle z_l^a z_{l'}^b \rangle = \delta_{l,l'} \frac{1}{RMP} \sum_{ss'} \left(\left(\sum_{\mu} u_{\mu s}^a u_{\mu s'}^b \right) \left(\sum_p v_{ps}^a v_{ps'}^b \right) \right) \quad (\text{B4})$$

$$= \frac{\delta_{l,l'}}{R} \sum_{ss'} \left(\left(\frac{1}{M} \sum_{\mu} u_{\mu s}^a u_{\mu s'}^b \right) \left(\frac{1}{P} \sum_p v_{ps}^a v_{ps'}^b \right) \right) \quad (\text{B5})$$

We now make the following assumption:

$$\frac{1}{M} \sum_{\mu} u_{\mu s}^a u_{\mu s'}^b = \begin{cases} Q_u^{ab} = O(1) & \text{if } s = s' \\ (Q_u^{ab})_{ss'} = O(\frac{1}{\sqrt{M}}) & \text{if } s \neq s' \end{cases} \quad (\text{B6})$$

This assumption corresponds to breaking the column-permutation symmetry and more generally the rotational symmetry between different replicas. We thus assume that the s -th column of \mathbf{U}^a is correlated to the s -th column of \mathbf{U}^b and to none of the others. We make the same assumption for \mathbf{V} . Then,

$$\langle z_l^a z_{l'}^b \rangle = \frac{\delta_{l,l'}}{R} \left(\sum_s Q_u^{ab} Q_v^{ab} + \sum_{s \neq s'} (Q_u^{ab})_{ss'} (Q_v^{ab})_{ss'} \right). \quad (\text{B7})$$

Due to the hypothesis (B6), the second term vanishes, and

$$\langle z_l^a z_{l'}^b \rangle = \delta_{s,s'} Q_u^{ab} Q_v^{ab}. \quad (\text{B8})$$

Note that by definition of Q_u^{ab} in (B6), $Q_u^{ab} = Q_u^{ba}$.

b. *Introducing $\hat{\mathbf{Q}}_u$* In equation (43), Dirac δ functions enforce the relations (B6). We use the integral representation of these δ functions to carry on the calculation:

$$\delta\left(MQ_u^{ab} - \sum_{\mu} u_{\mu s}^a u_{\mu s}^b\right) = \frac{1}{2\pi i} \int d\tilde{Q}_U^{ab} e^{-\tilde{Q}_U^{ab}(MQ_u^{ab} - \sum_{\mu} u_{\mu s}^a u_{\mu s}^b)}. \quad (\text{B9})$$

The product of all these δ functions thus gives

$$\prod_{a \leq b} \delta\left(MQ_u^{ab} - \sum_{\mu} u_{\mu s}^a u_{\mu s}^b\right) \propto \int d\tilde{\mathbf{Q}}_U e^{-M \sum_{a \leq b} \tilde{Q}_U^{ab} Q_u^{ab}} e^{\sum_{\mu} \sum_{a \leq b} \tilde{Q}_U^{ab} u_{\mu s}^a u_{\mu s}^b}. \quad (\text{B10})$$

Note that because $Q_u^{ab} = Q_u^{ba}$, the replica indices in the sum are $a \leq b$. Finally, we make a change of variables

$$\forall a, \quad \hat{Q}_U^{aa} = 2\tilde{Q}_U^{aa} \quad (\text{B11})$$

$$\forall(a, b) \text{ with } a \neq b, \quad \hat{Q}_U^{ab} = 4\tilde{Q}_U^{ab} \quad (\text{B12})$$

which allows us to obtain the following formulas

$$\sum_{a \leq b} \tilde{Q}_U^{ab} Q_u^{ab} = \frac{1}{2} \text{Tr}(\mathbf{Q}_u \hat{\mathbf{Q}}_u), \quad (\text{B13})$$

$$\sum_{a \leq b} \tilde{Q}_U^{ab} u_{\mu s}^a u_{\mu s}^b = \frac{1}{2} \mathbf{u}_{\mu s}^{\top} \hat{\mathbf{Q}}_u \mathbf{u}_{\mu s}, \quad (\text{B14})$$

where we introduced the vector $\mathbf{u}_{\mu s} = (u_{\mu s}^0 \dots u_{\mu s}^n)^{\top}$. We change the integration variable from $\tilde{\mathbf{Q}}_U$ to $\hat{\mathbf{Q}}_u$, and we obtain the expression (44).

-
- [1] D. B. Rubin A. P. Dempster, N. M. Laird. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.
 - [2] F. Caltagirone, L. Zdeborová, and F. Krzakala. On convergence of approximate message passing. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 1812–1816, June 2014.
 - [3] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
 - [4] David L Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.
 - [5] David L. Donoho, Matan Gavish, and Andrea Montanari. The phase transition of matrix recovery from Gaussian measurements matches the minimax mse of matrix denoising. *Proceedings of the National Academy of Sciences*, 110(21):8405–8410, 2013.
 - [6] David L Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
 - [7] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 665–674, New York, NY, USA, 2013. ACM.
 - [8] Y. Kabashima. An integral formula for large random rectangular matrices and its application to analysis of linear vector channels. In *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks and Workshops, 2008. WiOPT 2008. 6th International Symposium on*, pages 620–624, April 2008.
 - [9] Yoshiyuki Kabashima. A CDMA multiuser detection algorithm on the basis of belief propagation. *Journal of Physics A: Mathematical and General*, 36(43):11111, 2003.
 - [10] Yoshiyuki Kabashima, Florent Krzakala, Marc Mézard, Ayaka Sakata, and Lenka Zdeborová. Phase transitions and sample complexity in Bayes-optimal matrix factorization. *to appear in IEEE Trans. Inf. Theory, preprint arXiv:1402.1298*, 2016.
 - [11] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová. Statistical-physics-based reconstruction in compressed sensing. *Phys. Rev. X*, 2:021005, May 2012.
 - [12] Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Phase diagram and approximate message passing for blind calibration and dictionary learning. In *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pages 659–663. IEEE, 2013.
 - [13] Kiryung Lee, Yihong Wu, and Yoram Bresler. Near optimal compressed sensing of sparse rank-one matrices via sparse power factorization. *arXiv preprint arXiv:1312.0525*, 2013.
 - [14] T. Lesieur, F. Krzakala, and L. Zdeborová. Phase transitions in sparse PCA. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 1635–1639, June 2015.

- [15] Ryosuke Matsushita and Toshiyuki Tanaka. Low-rank matrix reconstruction and clustering via approximate message passing. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 917–925. Curran Associates, Inc., 2013.
- [16] Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [17] Marc Mézard, Giorgio Parisi, and Miguel-Angel Virasoro. *Spin glass theory and beyond*. 1990.
- [18] Hidetoshi Nishimori. *Statistical physics of spin glasses and information processing: an introduction*. Number 111. Oxford University Press, 2001.
- [19] Jason T Parker, Yan Shou, and Philip Schniter. Parametric bilinear generalized approximate message passing. *arXiv preprint arXiv:1508.07575*, 2015.
- [20] J.T. Parker, P. Schniter, and V. Cevher. Bilinear generalized approximate message passing x2014;Part ii: Applications. *Signal Processing, IEEE Transactions on*, 62(22):5854–5867, Nov 2014.
- [21] J.T. Parker, P. Schniter, and V. Cevher. Bilinear generalized approximate message passing;Part i: Derivation. *Signal Processing, IEEE Transactions on*, 62(22):5839–5853, Nov 2014.
- [22] S. Rangan. Generalized approximate message passing for estimation with random linear mixing. In *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pages 2168–2172, July 2011.
- [23] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [24] E. Riegler, D. Stotz, and H. Bolcskei. Information-theoretic limits of matrix completion. In *Information Theory (ISIT), 2015 IEEE International Symposium on*, pages 1836–1840, June 2015.
- [25] A. Sakata and Y. Kabashima. Sample complexity of Bayesian optimal dictionary learning. In *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pages 669–673, July 2013.
- [26] D. J. Thouless, P. W. Anderson, and R. G. Palmer. Solution of 'solvable model of a spin glass'. *Philosophical Magazine*, 35(3):593–601, 1977.
- [27] Jeremy Vila, Philip Schniter, Sundeep Rangan, Florent Krzakala, and Lenka Zdeborová. Adaptive damping and mean removal for the generalized approximate message passing algorithm. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 2021–2025. IEEE, 2015.
- [28] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *arXiv preprint arXiv:1511.02476*, 2015.
- [29] Qinqing Zheng and John Lafferty. A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 109–117. Curran Associates, Inc., 2015.